



US009323637B2

(12) **United States Patent**
Lucas et al.

(10) **Patent No.:** **US 9,323,637 B2**
(45) **Date of Patent:** **Apr. 26, 2016**

(54) **POWER SEQUENCING AND DATA
HARDENING ARCHITECTURE**

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,322,766 A	3/1982	Becker et al.
4,528,458 A	7/1985	Nelson et al.
4,600,962 A	7/1986	Bliehall

(Continued)

FOREIGN PATENT DOCUMENTS

EP	1465203 A1	10/2004
EP	1 956 489 A2	8/2008

(Continued)

OTHER PUBLICATIONS

International Search Report and Written Opinion dated May 27, 2015, received in International Patent Application No. PCT/US2014/067476, which corresponds to U.S. Appl. No. 14/135,417, 14 pages (Lucas).

(Continued)

Primary Examiner — Philip Guyton

(74) *Attorney, Agent, or Firm* — Morgan, Lewis & Bockius LLP

(71) Applicant: **SanDisk Enterprise IP LLC**, Milpitas, CA (US)

(72) Inventors: **Gregg S. Lucas**, Tucson, AZ (US);
Kenneth B. Delpapa, Natick, MA (US);
Lace J. Herman, Longmont, CO (US);
Robert W. Ellis, Phoenix, AZ (US)

(73) Assignee: **SANDISK ENTERPRISE IP LLC**, Milpitas, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 128 days.

(21) Appl. No.: **14/135,371**

(22) Filed: **Dec. 19, 2013**

(65) **Prior Publication Data**

US 2015/0100824 A1 Apr. 9, 2015

Related U.S. Application Data

(60) Provisional application No. 61/887,910, filed on Oct. 7, 2013.

(51) **Int. Cl.**
G06F 11/00 (2006.01)
G06F 11/30 (2006.01)
(Continued)

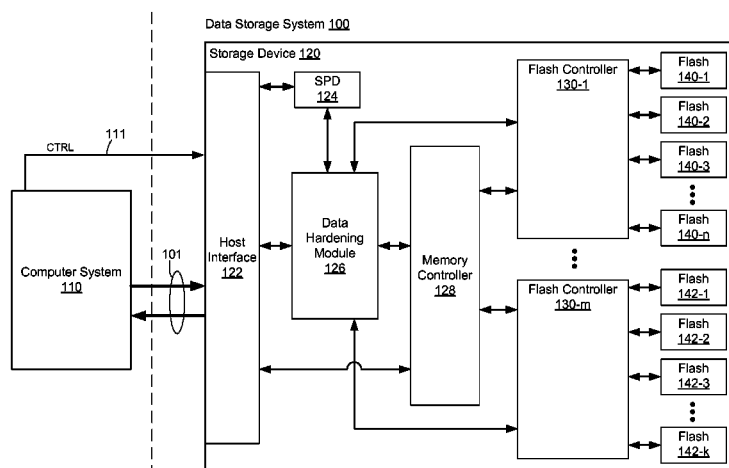
(52) **U.S. Cl.**
CPC **G06F 11/3058** (2013.01); **G06F 1/28**
(2013.01); **G06F 11/1402** (2013.01); **G06F**
11/1441 (2013.01)

(58) **Field of Classification Search**
CPC G06F 11/1412; G06F 11/1441; G06F
11/3058; G06F 1/28
See application file for complete search history.

(57) **ABSTRACT**

The various implementations described herein include systems, methods and/or devices used to enable power sequencing and a data hardening module in a storage device. In one aspect, the method includes determining whether a power supply voltage provided to the storage device is lower than an under-voltage threshold. The method further includes, in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, performing a power fail operation, the power fail operation including: (1) signaling a power fail condition to a plurality of controllers on the storage device, (2) transferring data held in volatile memory to non-volatile memory, and (3) removing power from the plurality of controllers on the storage device.

24 Claims, 10 Drawing Sheets



Page 2

- [illegible]

(56)

References Cited**U.S. PATENT DOCUMENTS**

2009/0003058	A1	1/2009	Kang	
2009/0037652	A1	2/2009	Yu et al.	
2009/0144598	A1	6/2009	Yoon et al.	
2009/0168525	A1	7/2009	Olbrich et al.	
2009/0172258	A1	7/2009	Olbrich et al.	
2009/0172259	A1	7/2009	Prins et al.	
2009/0172260	A1	7/2009	Olbrich et al.	
2009/0172261	A1	7/2009	Prins et al.	
2009/0172262	A1	7/2009	Olbrich et al.	
2009/0172308	A1	7/2009	Prins et al.	
2009/0172335	A1	7/2009	Kulkarni et al.	
2009/0172499	A1	7/2009	Olbrich et al.	
2009/0193058	A1	7/2009	Reid	
2009/0207660	A1	8/2009	Hwang et al.	
2009/0222708	A1	9/2009	Yamaga	
2009/0296466	A1	12/2009	Kim et al.	
2009/0296486	A1	12/2009	Kim et al.	
2009/0319864	A1	12/2009	Shrader	
2010/0008175	A1*	1/2010	Sweere et al.	365/229
2010/0011261	A1	1/2010	Cagno et al.	
2010/0052426	A1	3/2010	Carter et al.	
2010/0052625	A1	3/2010	Cagno et al.	
2010/0061151	A1	3/2010	Miwa et al.	
2010/0095048	A1	4/2010	Bechtolsheim et al.	
2010/0103737	A1	4/2010	Park	
2010/0199125	A1	8/2010	Reche	
2010/0202196	A1	8/2010	Lee et al.	
2010/0208521	A1	8/2010	Kim et al.	
2010/0217920	A1	8/2010	Song	
2010/0262889	A1	10/2010	Bains	
2010/0281207	A1	11/2010	Miller et al.	
2010/0281342	A1	11/2010	Chang et al.	
2011/0066872	A1*	3/2011	Miller et al.	713/340
2011/0083060	A1	4/2011	Sakurada et al.	
2011/0085657	A1	4/2011	Matthews, Jr.	
2011/0113281	A1	5/2011	Zhang et al.	
2011/0131444	A1	6/2011	Buch et al.	
2011/0205823	A1	8/2011	Hemink et al.	
2011/0213920	A1	9/2011	Frost et al.	
2011/0228601	A1	9/2011	Olbrich et al.	
2011/0231600	A1	9/2011	Tanaka et al.	
2012/0054456	A1	3/2012	Grube et al.	
2012/0084492	A1	4/2012	Stenfort	
2012/0089855	A1	4/2012	Beckhoff et al.	
2012/0195126	A1	8/2012	Roohparvar	
2012/0221801	A1	8/2012	Okawa	
2012/0239976	A1	9/2012	Cometti et al.	
2012/0271990	A1*	10/2012	Chen et al.	711/103
2013/0019076	A1*	1/2013	Amidi et al.	711/162
2013/0336081	A1	12/2013	Sheets et al.	
2014/0001861	A1	1/2014	Mann et al.	
2014/0006798	A1	1/2014	Prakash et al.	
2014/0008970	A1	1/2014	Yamaguchi	
2014/0012522	A1	1/2014	Colombi et al.	
2014/0215103	A1	7/2014	Cohen et al.	
2014/0269053	A1	9/2014	Chen et al.	
2015/0052397	A1*	2/2015	Nakamura et al.	714/22

FOREIGN PATENT DOCUMENTS

JP	2002-532806	S	10/2002
WO	WO 2007/036834	A2	4/2007
WO	WO 2007/080586	A2	7/2007
WO	WO 2008/121553	A1	10/2008
WO	WO 2008/121577	A1	10/2008
WO	WO 2009/028281	A1	3/2009
WO	WO 2009/032945	A1	3/2009
WO	WO 2009/058140	A1	5/2009
WO	WO 2009/134576	A1	11/2009

OTHER PUBLICATIONS

IBM Corporation, "Systems Management, Work Management," Version 5, Release 4, 9th Edition, Feb. 2006, pp. 1-21.

Texas instruments, "Power Management IC for Digital Set Top Boxes," SLVSA10A, Sep. 2009, pp. 1-22.

International Search Report and Written Opinion dated Jan. 26, 2015, received in International Patent Application No. PCT/US2014/059118, which corresponds to U.S. Appl. No. 14/135,371, 11 pages (Lucas).

International Search Report and Written, Opinion dated Jul. 26, 2013, received in International Patent Application No. PCT/US2013/035162, which corresponds to U.S. Appl. No. 13/855,567, 7 pages (Ellis).

International Preliminary Report on Patentability dated Oct. 30, 2014, received in International Patent Application No. PCT/US2013/035162, which corresponds to U.S. Appl. No. 13/855,567, 4 pages (Ellis).

Barr, Introduction to Watchdog Timers, Oct. 2001, 3 pgs.

Canim, Buffered Bloom filters on Solid State Storage, ADMS*10, Singapore, Sep. 13-17, 2010, 8 pgs.

Kang, A Multi-Channel Architecture for High-Performance NAND Flash-Based Storage System, J. Syst. Archit., 53, 9, Sep. 2007, 15 pgs.

Kim, A Space-Efficient Flash Translation Layer for CompactFlash Systems, May 2002, 10 pgs.

Lu, A Forest-structured Bloom Filter with Flash Memory, MSST 2011, Denver, CO, May 23-27, 2011, article, 6 pgs.

Lu, A Forest-structured Bloom Filter with Flash Memory, MSST 2011, Denver, CO, May 23-27, 2011, presentation slides, 25 pgs.

McLean, Information Technology-AT Attachment with Packet Interface Extension, Aug. 19, 1998, 339 pgs.

Park, A High Performance Controller for NAND Flash-Based Solid State Disk (NSSD), Feb. 12-16, 2006, 4 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88133, Mar. 19, 2009, 7 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88136, Mar. 19, 2009, 7 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88146, Feb. 26, 2009, 10 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88154, Feb. 27, 2009, 8 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88164, Feb. 13, 2009, 6 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88206, Feb. 18, 2009, 8 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88217, Feb. 19, 2009, 7 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88229, Feb. 13, 2009, 7 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88232, Feb. 19, 2009, 8 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US08/88236, Feb. 19, 2009, 7 pgs.

Pliant Technology, International Search Report / Written Opinion, PCT/US2011/028637, Oct. 27, 2011, 11 pgs.

Pliant Technology, Supplementary ESR, 08866997.3, Feb. 23, 2012, 6 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/042764, Aug. 31, 2012, 12 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/042771, Mar. 4, 2013, 14 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/042775, Sep. 26, 2012, 8 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/059447, Jun. 6, 2013, 12 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/059453, Jun. 6, 2013, 12 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/059459, Feb. 14, 2013, 9 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/065914, May 23, 2013, 7 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/065916, Apr. 5, 2013, 7 pgs.

SanDisk Enterprise IP LLC, International Search Report / Written Opinion, PCT/US2012/065919, Jun. 17, 2013, 8 pgs.

(56)

References Cited

OTHER PUBLICATIONS

SanDisk Enterprise IP LLC, Notification of the Decision to Grant a Patent Right for Patent for Invention, CN 200880127623.8, Jul. 4, 2013, 1 pg.

SanDisk Enterprise IP LLC, Office Action, CN 200880127623.8, Apr. 18, 2012, 12 pgs.

SanDisk Enterprise IP LLC, Office Action, CN 200880127623.8, Dec. 31, 2012, 9 pgs.

SanDisk Enterprise IP LLC, Office Action, JP 2010-540863, Jul. 24, 2012, 3 pgs.

Watchdog Timer and Power Savin Modes, Microchip Technology Inc., 2005, 14 pgs.

Zeidman, 1999 Verilog Designer's Library, 9 pgs.

International Search Report and Written Opinion dated Jul. 14, 2015, received in International Patent Application No. PCT/US2015/027263, which corresponds to U.S. Appl. No. 14/599,128, 10 pages (Ellis).

* cited by examiner

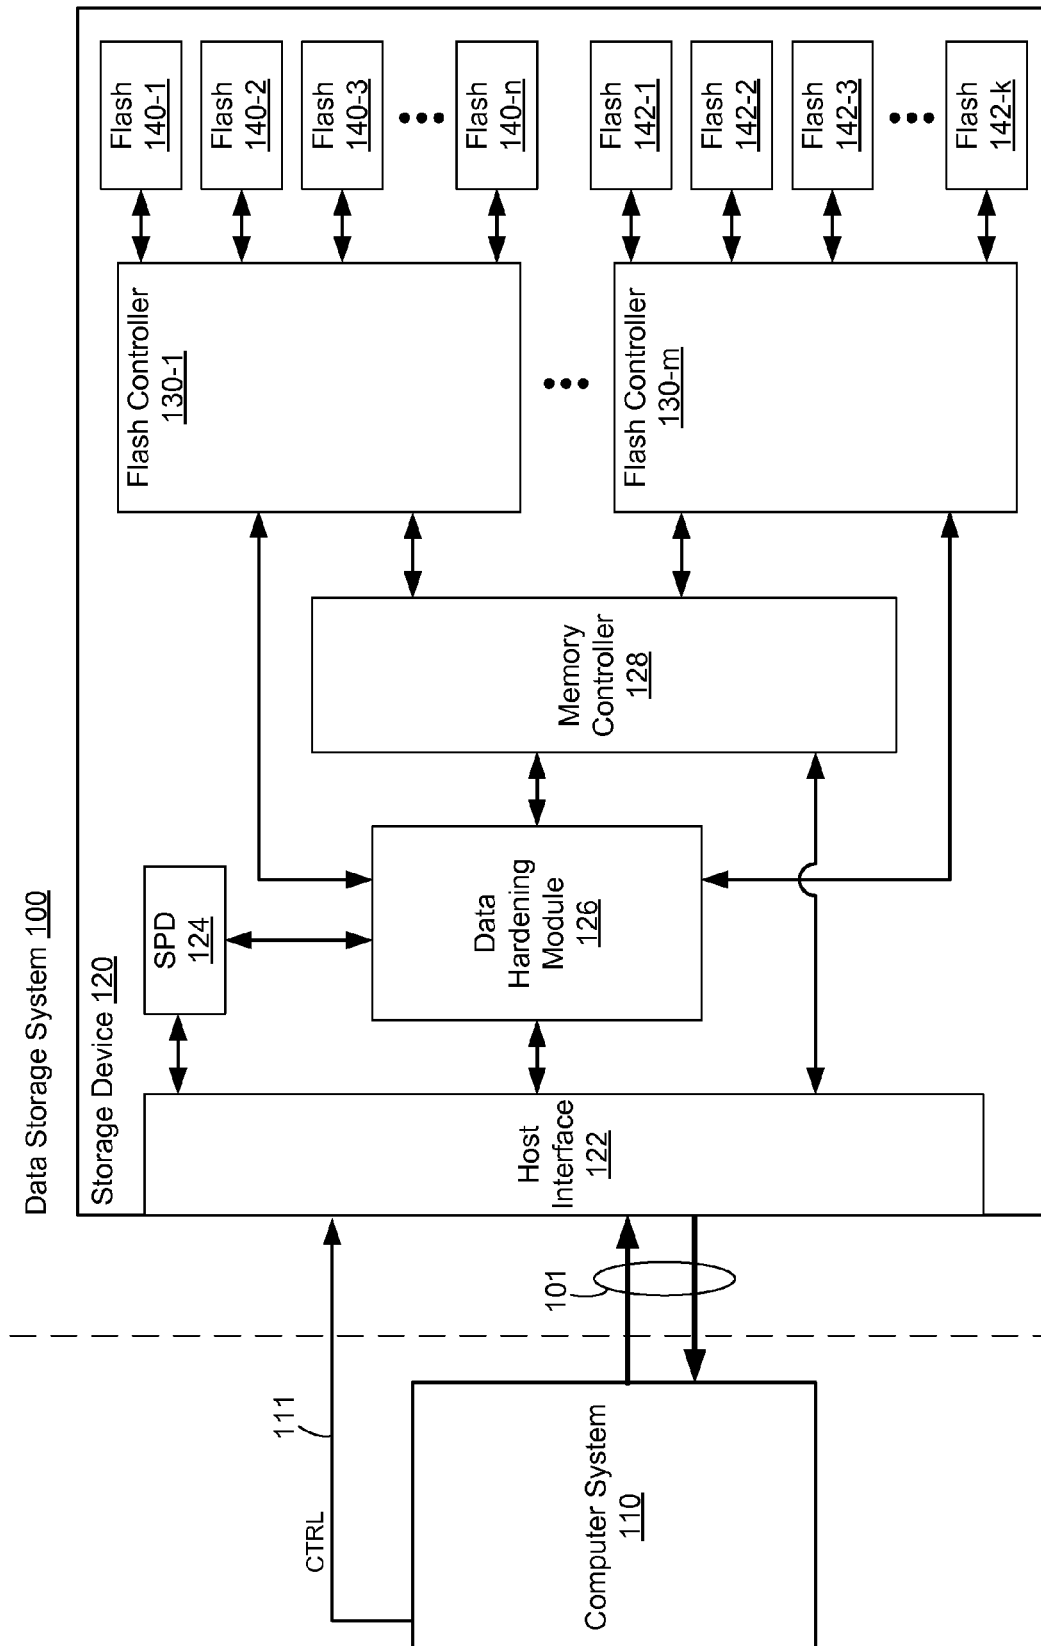


Figure 1

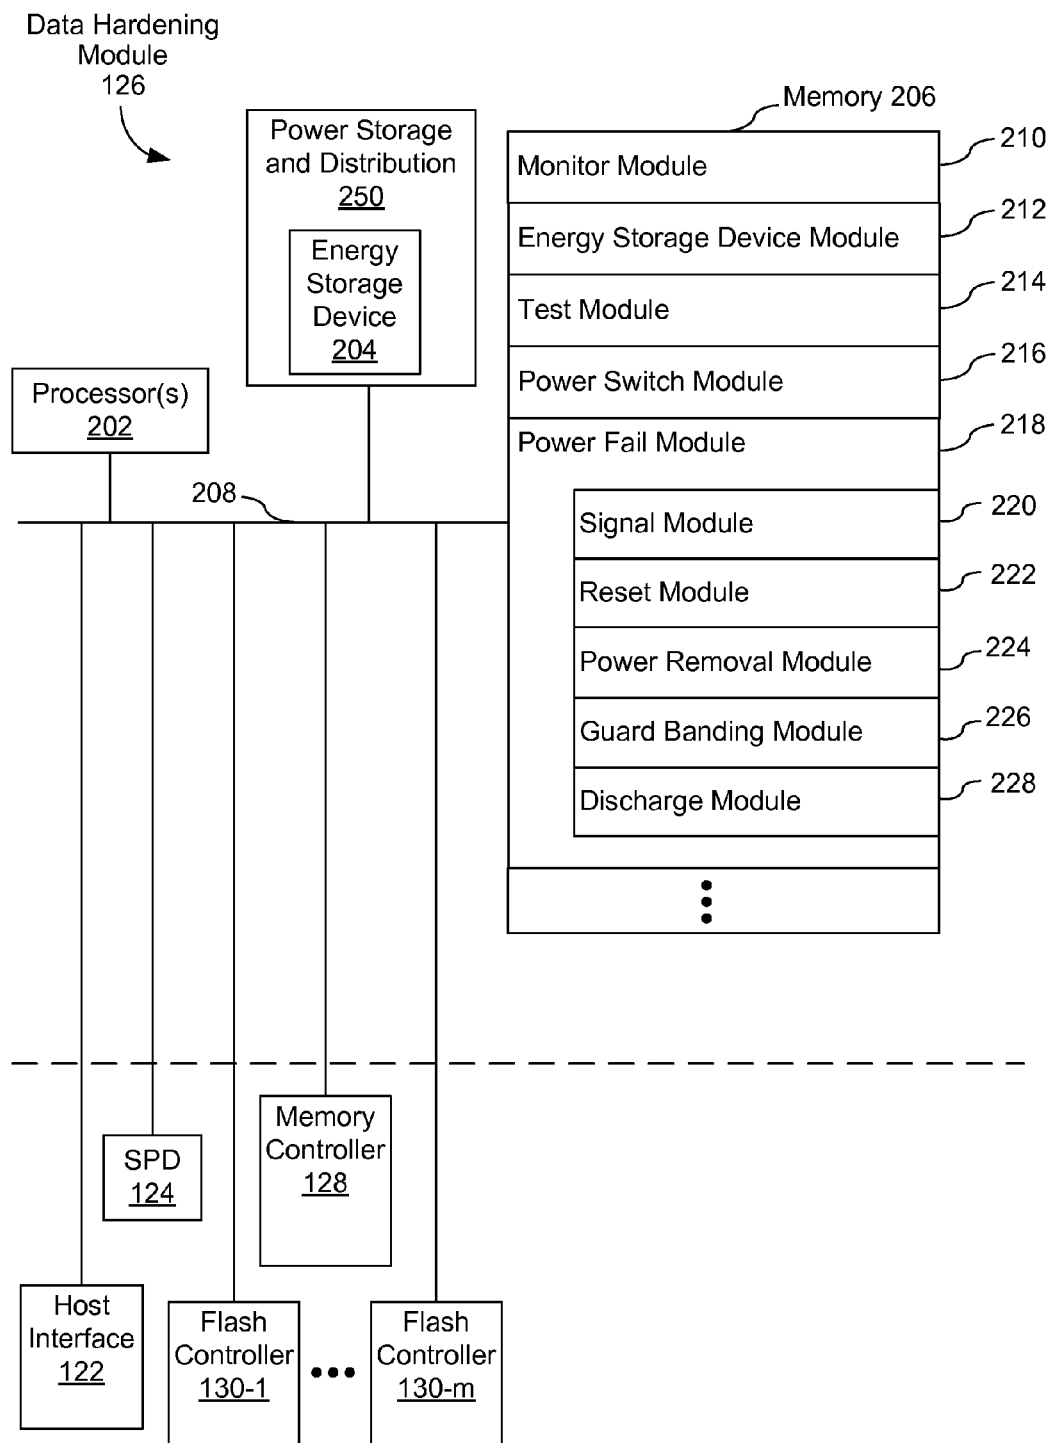


Figure 2

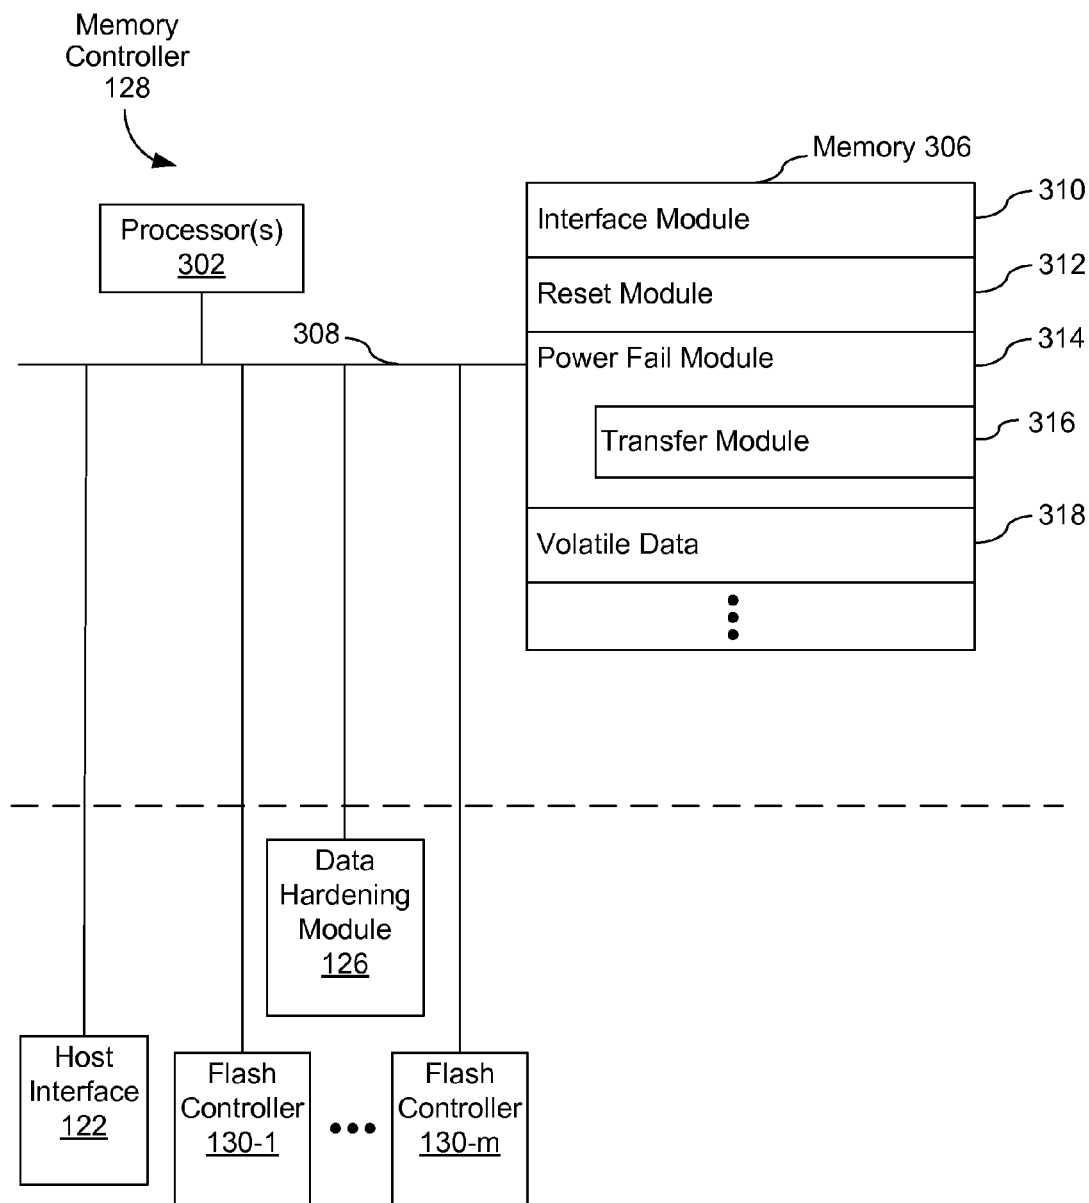


Figure 3

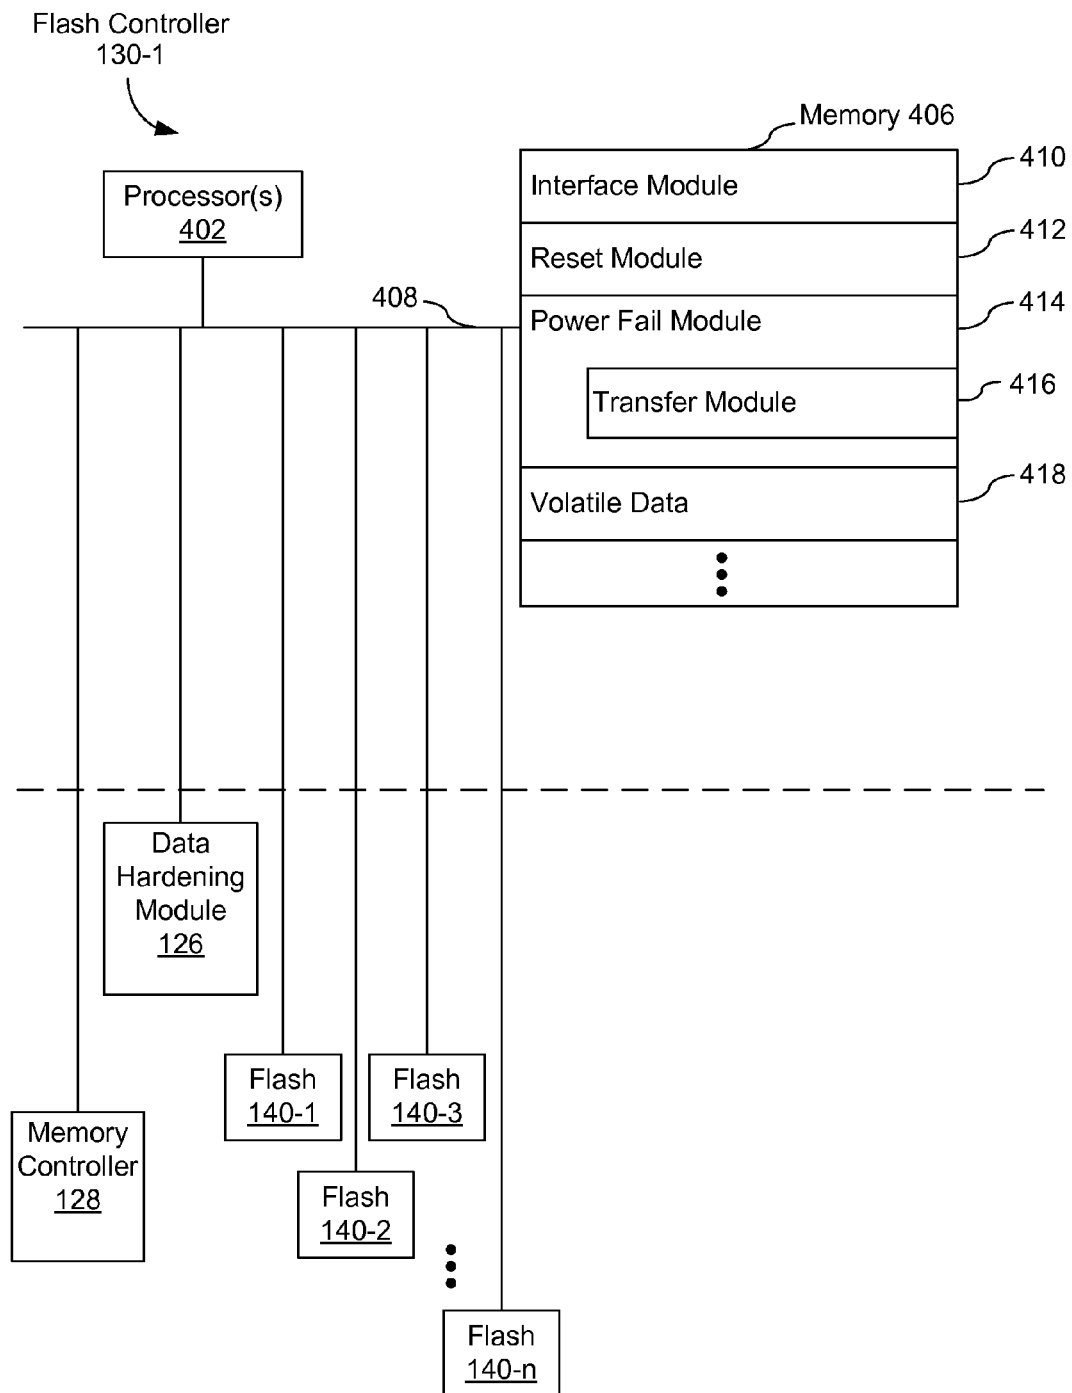


Figure 4

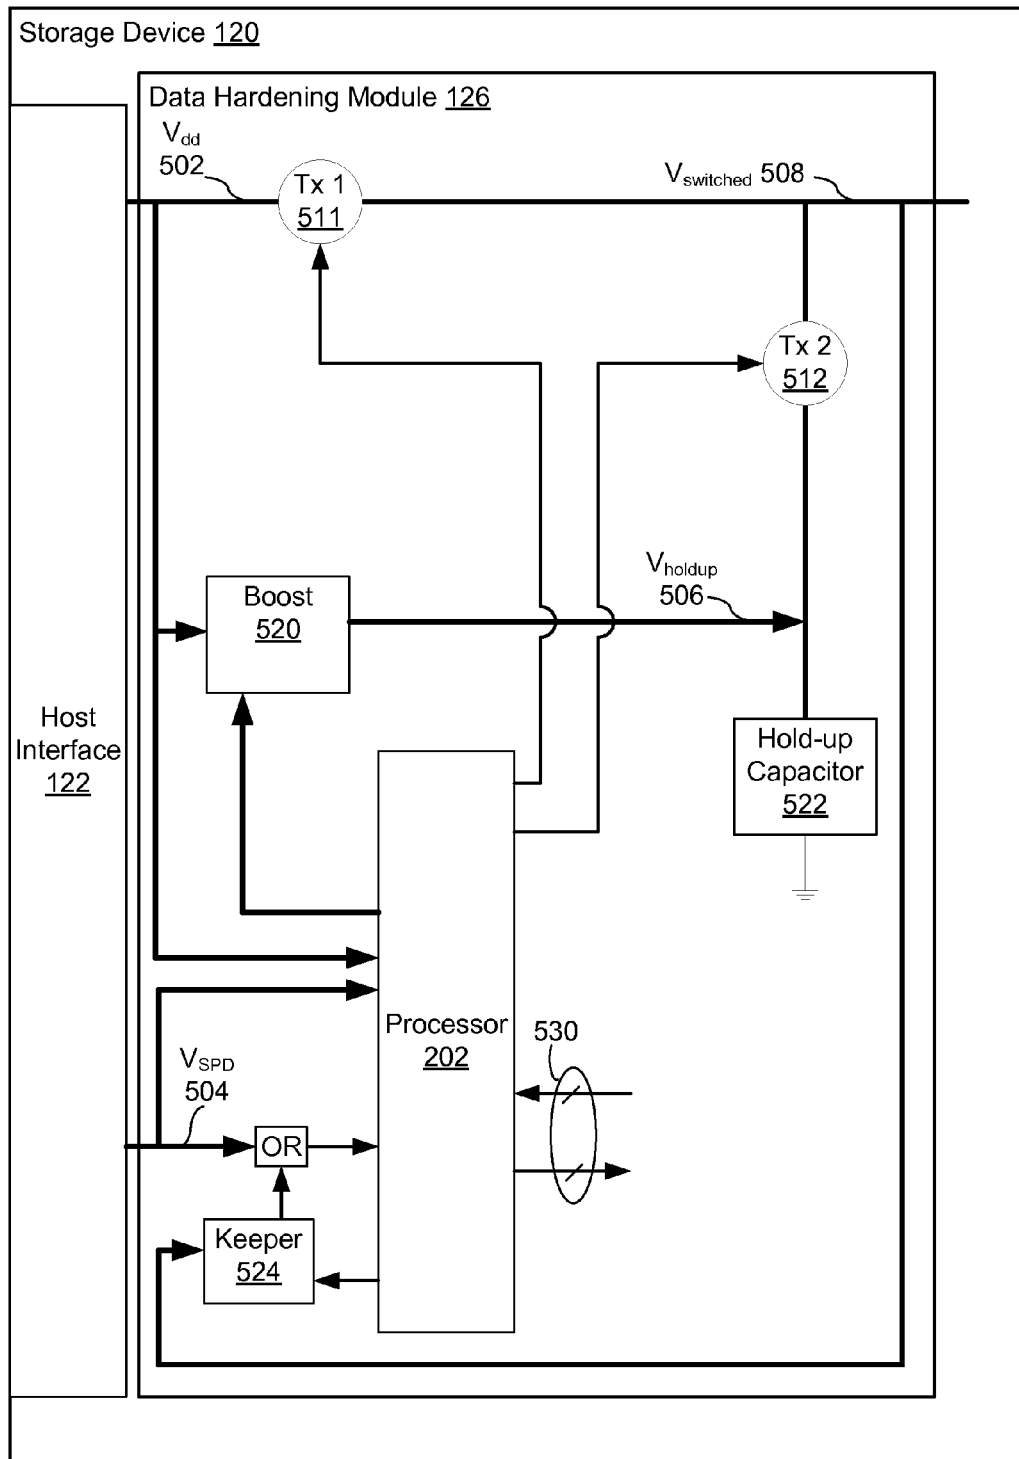


Figure 5

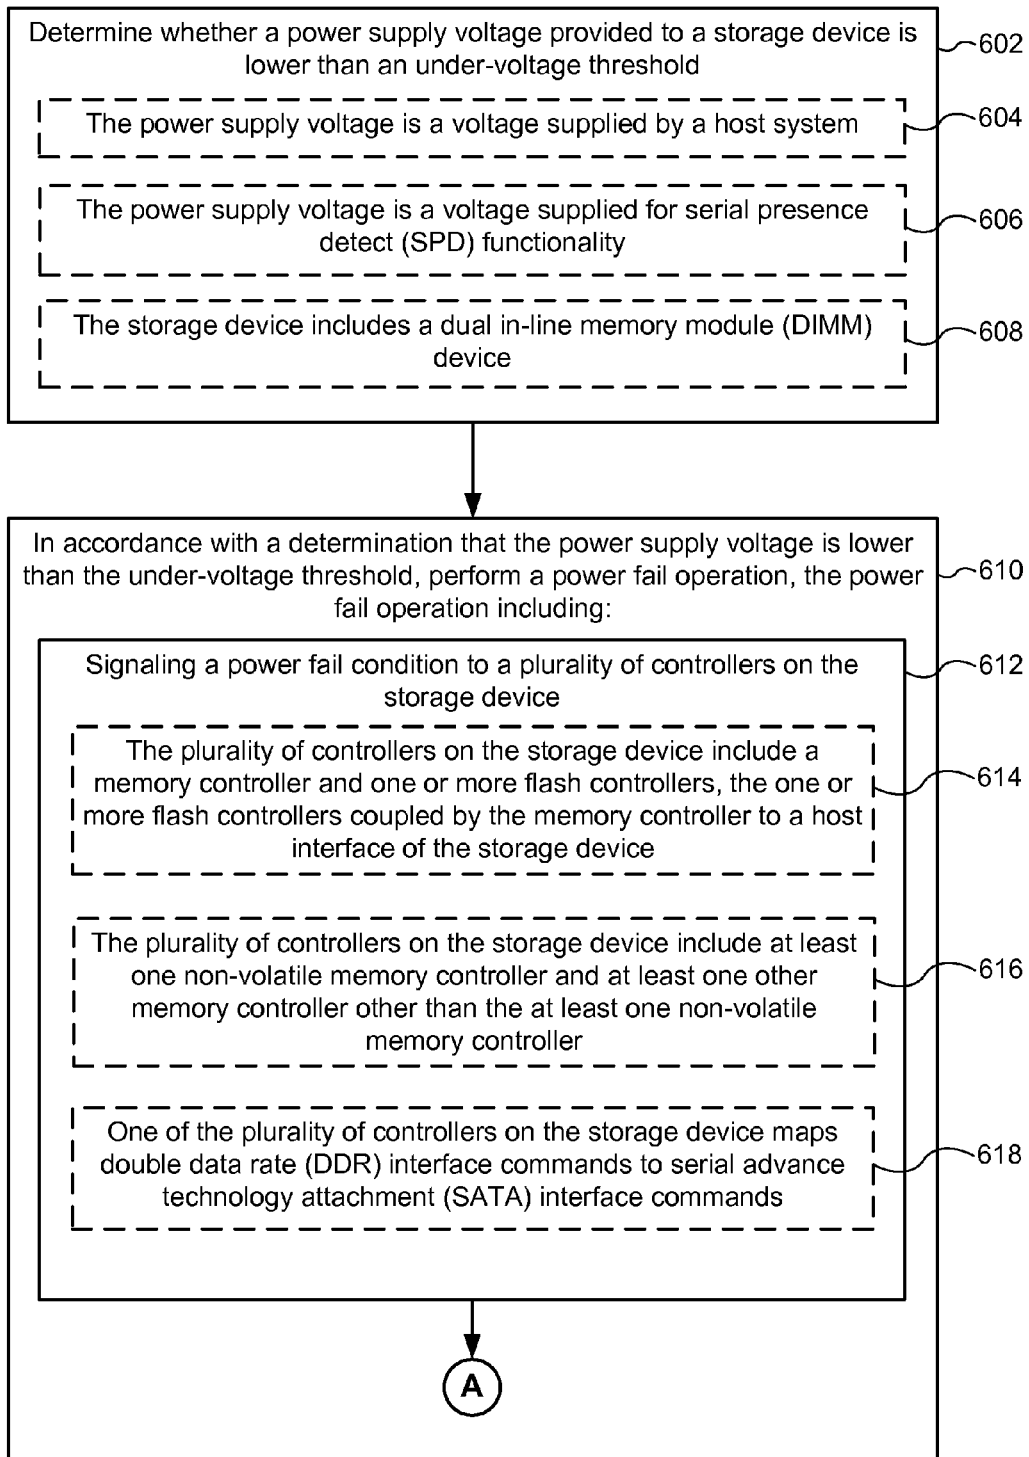
600

Figure 6A

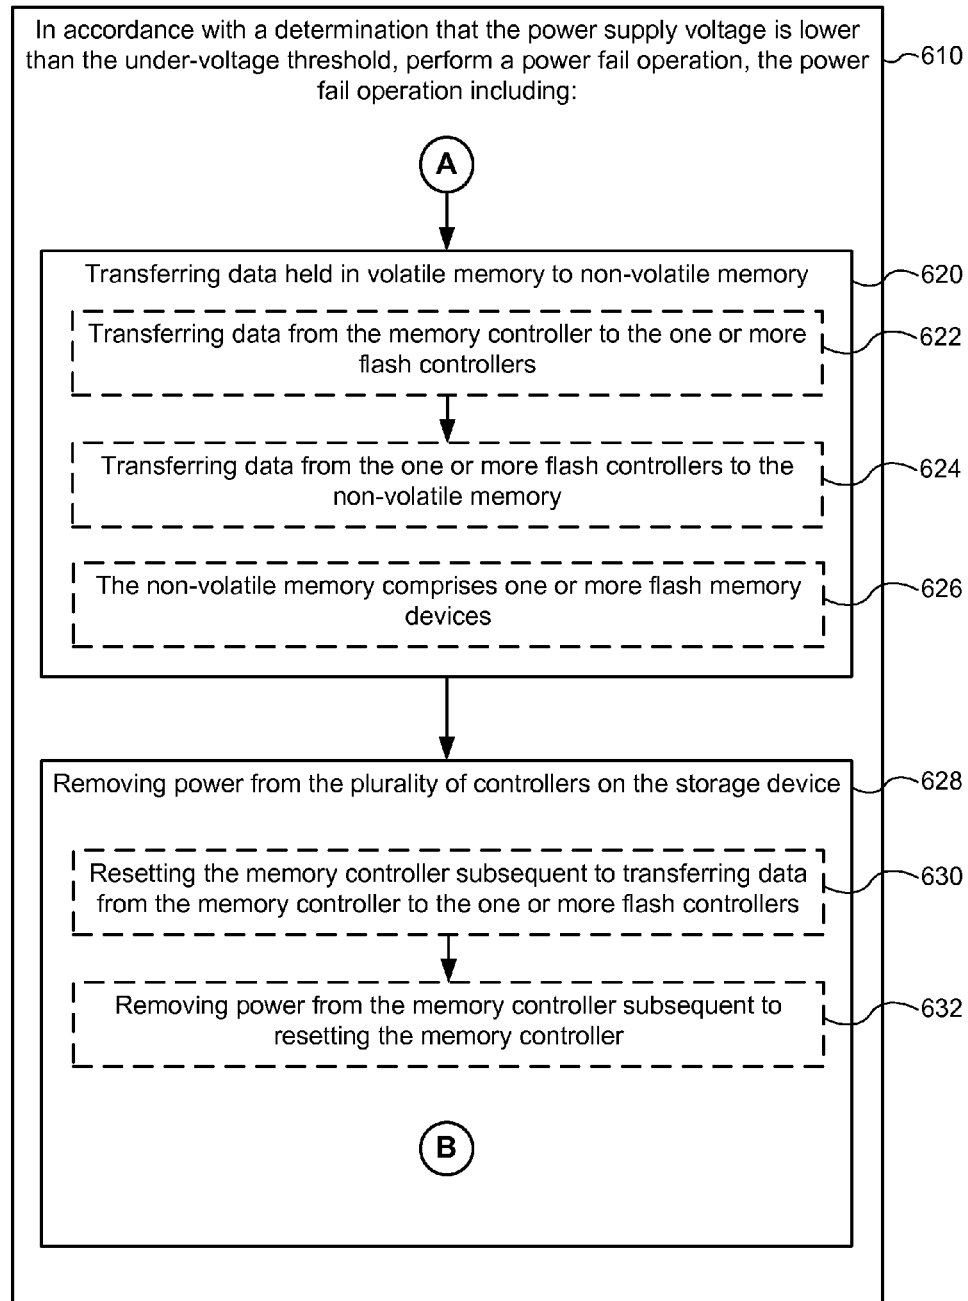


Figure 6B

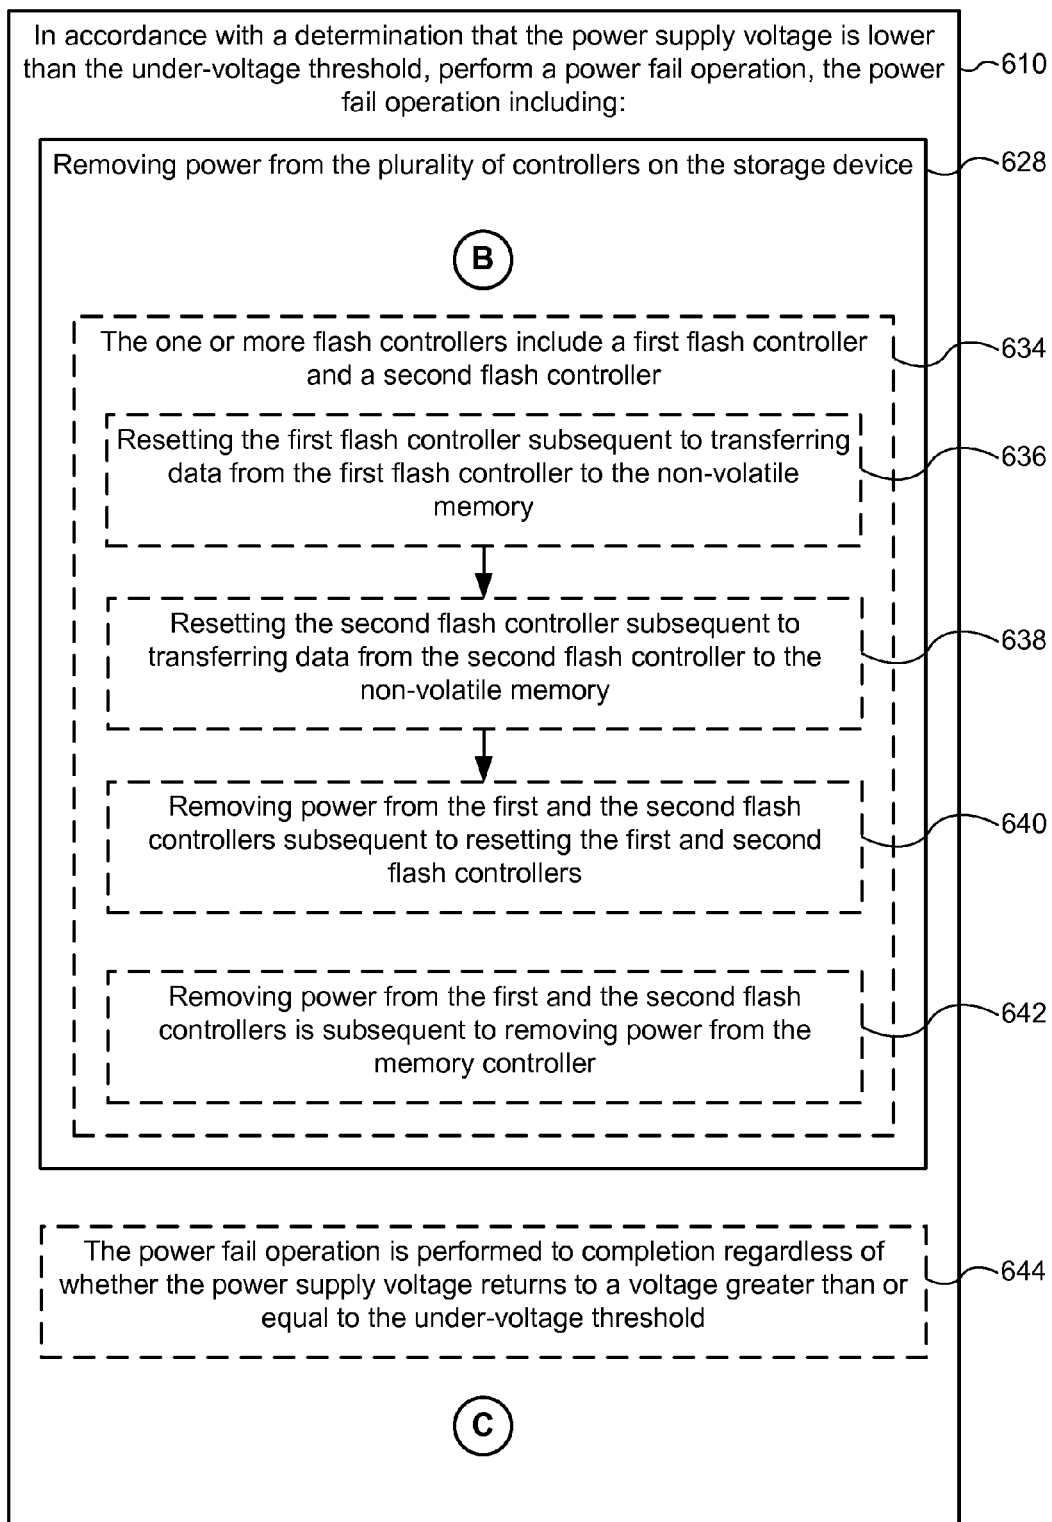


Figure 6C

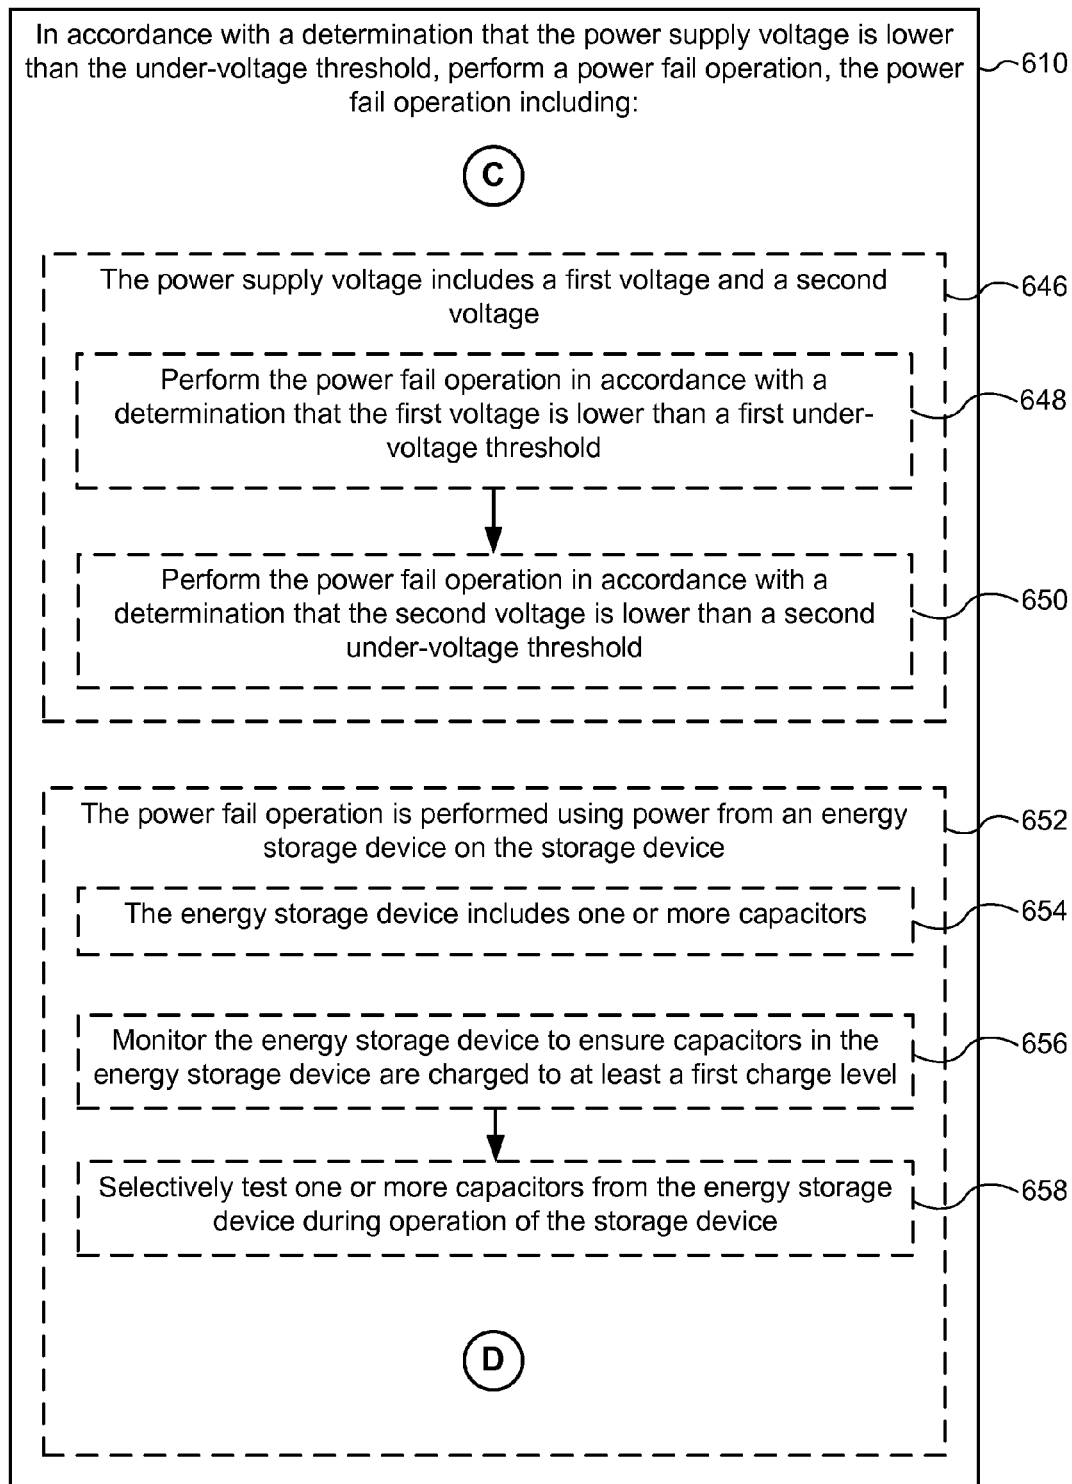


Figure 6D

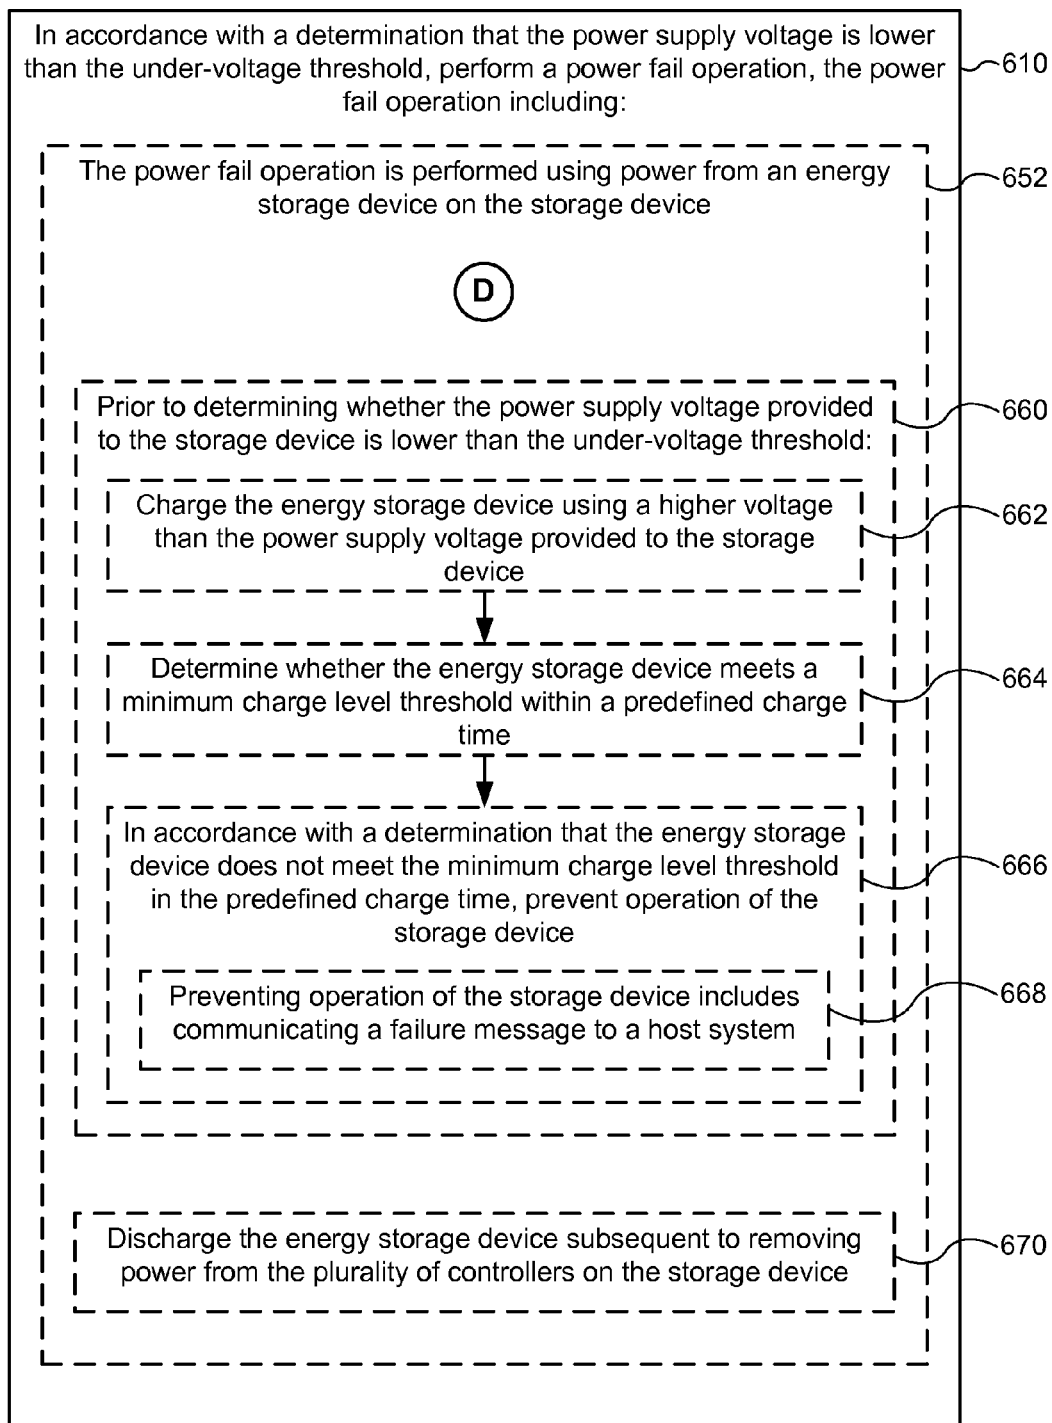


Figure 6E

1

POWER SEQUENCING AND DATA HARDENING ARCHITECTURE

RELATED APPLICATIONS

This application claims priority to U.S. Provisional Patent Application 61/887,910, filed Oct. 7, 2013, which is hereby incorporated by reference in its entirety.

TECHNICAL FIELD

The disclosed embodiments relate generally to memory systems, and in particular, to power sequencing and data hardening architecture in memory devices.

BACKGROUND

Semiconductor memory devices, including flash memory, typically utilize memory cells to store data as an electrical value, such as an electrical charge or voltage. A flash memory cell, for example, includes a single transistor with a floating gate that is used to store a charge representative of a data value. Flash memory is a non-volatile data storage device that can be electrically erased and reprogrammed. More generally, non-volatile memory (e.g., flash memory, as well as other types of non-volatile memory implemented using any of a variety of technologies) retains stored information even when not powered, as opposed to volatile memory, which requires power to maintain the stored information.

Data hardening, the saving of data and mission critical metadata held in volatile storage, is an integral part of a storage device. When there is a power failure, mission critical data may reside in volatile memory in a number of sub-system components. Coordinating and managing multiple sub-system components to ensure that volatile data is saved successfully is important for safeguarding data integrity of a storage device.

SUMMARY

Various implementations of systems, methods and devices within the scope of the appended claims each have several aspects, no single one of which is solely responsible for the attributes described herein. Without limiting the scope of the appended claims, after considering this disclosure, and particularly after considering the section entitled "Detailed Description" one will understand how the aspects of various implementations are used to enable power sequencing and data hardening in memory devices. In one aspect, a power fail operation is performed in accordance with a determination that a power supply voltage provided to a storage device is lower than an under-voltage threshold.

BRIEF DESCRIPTION OF THE DRAWINGS

So that the present disclosure can be understood in greater detail, a more particular description may be had by reference to the features of various implementations, some of which are illustrated in the appended drawings. The appended drawings, however, merely illustrate the more pertinent features of the present disclosure and are therefore not to be considered limiting, for the description may admit to other effective features.

FIG. 1 is a block diagram illustrating an implementation of a data storage system, in accordance with some embodiments.

2

FIG. 2 is a block diagram illustrating an implementation of a data hardening module, in accordance with some embodiments.

FIG. 3 is a block diagram illustrating an implementation of a memory controller, in accordance with some embodiments.

FIG. 4 is a block diagram illustrating an implementation of a flash controller, in accordance with some embodiments.

FIG. 5 is a block diagram illustrating an implementation of a data hardening module, in accordance with some embodiments.

FIGS. 6A-6E illustrate a flowchart representation of a method of protecting data in a storage device, in accordance with some embodiments.

In accordance with common practice the various features illustrated in the drawings may not be drawn to scale. Accordingly, the dimensions of the various features may be arbitrarily expanded or reduced for clarity. In addition, some of the drawings may not depict all of the components of a given system, method or device. Finally, like reference numerals may be used to denote like features throughout the specification and figures.

DETAILED DESCRIPTION

The various implementations described herein include systems, methods and/or devices used to enable power sequencing and data hardening in memory devices. Some implementations include systems, methods and/or devices to perform a power fail operation in accordance with a determination that a power supply voltage provided to a storage device is lower than an under-voltage threshold.

More specifically, some implementations include a method of protecting data in a storage device. In some implementations, the method includes determining whether a power supply voltage provided to the storage device is lower than an under-voltage threshold. The method further includes, in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, performing a power fail operation, the power fail operation including: (1) signaling a power fail condition to a plurality of controllers on the storage device, (2) transferring data held in volatile memory to non-volatile memory, and (3) removing power from the plurality of controllers on the storage device.

In some embodiments, the plurality of controllers on the storage device include a memory controller and one or more flash controllers, the one or more flash controllers coupled by the memory controller to a host interface of the storage device.

In some embodiments, transferring data held in volatile memory to non-volatile memory includes: (1) transferring data from the memory controller to the one or more flash controllers, and (2) transferring data from the one or more flash controllers to the non-volatile memory.

In some embodiments, removing power from the plurality of controllers on the storage device includes: (1) resetting the memory controller subsequent to transferring data from the memory controller to the one or more flash controllers, and (2) removing power from the memory controller subsequent to resetting the memory controller.

In some embodiments, the one or more flash controllers include a first flash controller and a second flash controller, and removing power from the plurality of controllers on the storage device includes: (1) resetting the first flash controller subsequent to transferring data from the first flash controller to the non-volatile memory, (2) resetting the second flash controller subsequent to transferring data from the second flash controller to the non-volatile memory, and (3) removing

power from the first and the second flash controllers subsequent to resetting the first and second flash controllers.

In some embodiments, removing power from the first and the second flash controllers is subsequent to removing power from the memory controller.

In some embodiments, the power fail operation is performed to completion regardless of whether the power supply voltage returns to a voltage greater than or equal to the under-voltage threshold.

In some embodiments, the power supply voltage is a voltage supplied by a host system.

In some embodiments, the power supply voltage is a voltage supplied for serial presence detect (SPD) functionality.

In some embodiments, the power supply voltage includes a first voltage and a second voltage, and performing the power fail operation includes: (1) performing the power fail operation in accordance with a determination that the first voltage is lower than a first under-voltage threshold, and (2) performing the power fail operation in accordance with a determination that the second voltage is lower than a second under-voltage threshold.

In some embodiments, the power fail operation is performed using power from an energy storage device on the storage device.

In some embodiments, the energy storage device includes one or more capacitors.

In some embodiments, the method further includes (1) monitoring the energy storage device to ensure capacitors in the energy storage device are charged to at least a first charge level, and (2) selectively testing one or more capacitors from the energy storage device during operation of the storage device.

In some embodiments, the method further includes, prior to determining whether the power supply voltage provided to the storage device is lower than the under-voltage threshold, (1) charging the energy storage device using a higher voltage than the power supply voltage provided to the storage device, (2) determining whether the energy storage device meets a minimum charge level threshold within a predefined charge time, and (3) in accordance with a determination that the energy storage device does not meet the minimum charge level threshold in the predefined charge time, preventing operation of the storage device.

In some embodiments, preventing operation of the storage device includes communicating a failure message to a host system.

In some embodiments, the method further includes discharging the energy storage device subsequent to removing power from the plurality of controllers on the storage device.

In some embodiments, the non-volatile memory comprises one or more flash memory devices.

In some embodiments, the storage device includes a dual in-line memory module (DIMM) device.

In some embodiments, the plurality of controllers on the storage device include at least one non-volatile memory controller and at least one other memory controller other than the at least one non-volatile memory controller.

In some embodiments, one of the plurality of controllers on the storage device maps double data rate (DDR) interface commands to serial advance technology attachment (SATA) interface commands.

In another aspect, any of the methods described above are performed by a storage device including (1) an interface for coupling the storage device to a host system, (2) a plurality of controllers, each of the plurality of controllers configured to transfer data held in volatile memory to non-volatile memory, and (3) a data hardening module including one or more pro-

cessors and an energy storage device, the data hardening module configured to: (a) determine whether a power supply voltage provided to the storage device is lower than an under-voltage threshold, and (b) in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, perform a power fail operation, the power fail operation including: (i) signaling a power fail condition to the plurality of controllers, causing the plurality of controllers to transfer data held in volatile memory to non-volatile memory, and (ii) removing power from the plurality of controllers on the storage device.

In some embodiments, the storage device is configured to perform any of the methods described above.

In yet another aspect, any of the methods described above are performed by a storage device operable to protect data. In some embodiments, the device includes (1) an interface for coupling the storage device to a host system, (2) means for determining whether a power supply voltage provided to the storage device is lower than an under-voltage threshold, and (3) means for performing a power fail operation, in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, the means for performing the power fail operation including: (a) means for signaling a power fail condition to a plurality of controllers on the storage device, (b) means for transferring data held in volatile memory to non-volatile memory, and (c) means for removing power from the plurality of controllers on the storage device.

In yet another aspect, a non-transitory computer readable storage medium, storing one or more programs for execution by one or more processors of a storage device having a plurality of controllers and a data hardening module, the one or more programs including instructions for performing any of the methods described above.

In some embodiments, the non-transitory computer readable storage medium includes a non-transitory computer readable storage medium associated with each of the plurality of controllers on the storage device and a non-transitory computer readable storage medium associated with the data hardening module.

Numerous details are described herein in order to provide a thorough understanding of the example implementations illustrated in the accompanying drawings. However, some embodiments may be practiced without many of the specific details, and the scope of the claims is only limited by those features and aspects specifically recited in the claims. Furthermore, well-known methods, components, and circuits have not been described in exhaustive detail so as not to unnecessarily obscure more pertinent aspects of the implementations described herein.

FIG. 1 is a block diagram illustrating an implementation of a data storage system 100, in accordance with some embodiments. While some example features are illustrated, various other features have not been illustrated for the sake of brevity and so as not to obscure more pertinent aspects of the example implementations disclosed herein. To that end, as a non-limiting example, data storage system 100 includes storage device 120, which includes host interface 122, serial presence detect (SPD) device 124, data hardening module 126, memory controller 128, one or more flash controllers (e.g., flash controller(s) 130), and non-volatile memory (e.g., one or more flash memory device(s) 140, 142), and is used in conjunction with computer system 110. In some implementations, storage device 120 (sometimes herein called a data storage device, information storage device, dual in-line memory module, or memory device) includes a single flash memory device while in other implementations storage

device **120** includes a plurality of flash memory devices. In some implementations, flash memory devices **140**, **142** include NAND-type flash memory or NOR-type flash memory. Further, in some implementations, flash controller **130** is a solid-state drive (SSD) controller. However, one or more other types of storage media may be included in accordance with aspects of a wide variety of implementations.

Computer system **110** is coupled to storage device **120** through data connections **101**. However, in some implementations computer system **110** includes storage device **120** as a component and/or sub-system. Computer system **110** may be any suitable computer device, such as a personal computer, a workstation, a computer server, or any other computing device. Computer system **110** is sometimes called a host or host system. In some implementations, computer system **110** includes one or more processors, one or more types of memory, optionally includes a display and/or other user interface components such as a keyboard, a touch screen display, a mouse, a track-pad, a digital camera and/or any number of supplemental devices to add functionality. Further, in some implementations, computer system **110** sends one or more host commands (e.g., read commands and/or write commands) on control line **111** to storage device **120**. In some implementations, computer system **110** is a server system, such as a server system in a data center, and does not have a display and other user interface components.

In some implementations, storage device **120** includes flash memory devices **140**, **142** (e.g., flash memory devices **140-1** through **140-n** and flash memory devices **142-1** through **142-k**) and flash controllers **130** (e.g., flash controllers **130-1** through **130-m**). In some implementations, each flash controller of flash controllers **130** include one or more processing units (also sometimes called CPUs or processors or microprocessors or microcontrollers) configured to execute instructions in one or more programs (e.g., in flash controllers **130**). In some implementations, the one or more processors are shared by one or more components within, and in some cases, beyond the function of flash controllers **130**. Flash memory devices **140**, **142** are coupled to flash controllers **130** through connections that typically convey commands in addition to data, and optionally convey metadata, error correction information and/or other information in addition to data values to be stored in flash memory devices **140**, **142** and data values read from flash memory devices **140**, **142**. For example, flash memory devices **140**, **142** can be configured for enterprise storage suitable for applications such as cloud computing, or for caching data stored (or to be stored) in secondary storage, such as hard disk drives. Additionally and/or alternatively, flash memory can also be configured for relatively smaller-scale applications such as personal flash drives or hard-disk replacements for personal, laptop and tablet computers. Although flash memory devices and flash controllers are used as an example here, storage device **120** may include any other non-volatile memory device(s) and corresponding non-volatile memory controller(s).

In some implementations, storage device **120** also includes host interface **122**, SPD device **124**, data hardening module **126**, and memory controller **128**. Storage device **120** may include various additional features that have not been illustrated for the sake of brevity and so as not to obscure more pertinent features of the example implementations disclosed herein, and a different arrangement of features may be possible. Host interface **122** provides an interface to computer system **110** through data connections **101**.

In some implementations, data hardening module **126** includes one or more processing units (also sometimes called CPUs or processors or microprocessors or microcontrollers)

configured to execute instructions in one or more programs (e.g., in data hardening module **126**). In some implementations, the one or more processors are shared by one or more other components (e.g., memory controller **128**) of storage device **120**. Data hardening module **126** is coupled to host interface **122**, SPD device **124**, memory controller **128**, and flash controllers **130** in order to coordinate the operation of these components, including supervising and controlling functions such as power up, power down, data hardening, charging energy storage device(s), data logging, and other aspects of managing functions on storage device **120**.

Memory controller **128** is coupled to host interface **122**, data hardening module **126**, and flash controllers **130**. In some implementations, during a write operation, memory controller **128** receives data from computer system **110** through host interface **122** and during a read operation, memory controller **128** sends data to computer system **110** through host interface **122**. Further, host interface **122** provides additional data, signals, voltages, and/or other information needed for communication between memory controller **128** and computer system **110**. In some embodiments, memory controller **128** and host interface **122** use a defined interface standard for communication, such as double data rate type three synchronous dynamic random access memory (DDR3). In some embodiments, memory controller **128** and flash controllers **130** use a defined interface standard for communication, such as serial advance technology attachment (SATA). In some other implementations, the device interface used by memory controller **128** to communicate with flash controllers **130** is SAS (serial attached SCSI), or other storage interface. In some implementations, memory controller **128** includes one or more processing units (also sometimes called CPUs or processors or microprocessors or microcontrollers) configured to execute instructions in one or more programs (e.g., in memory controller **128**). In some implementations, the one or more processors are shared by one or more components within, and in some cases, beyond the function of memory controller **128**.

SPD device **124** is coupled to host interface **122** and data hardening module **126**. Serial presence detect (SPD) refers to a standardized way to automatically access information about a computer memory module (e.g., storage device **120**). For example, if the memory module has a failure, the failure can be communicated with a host system (e.g., computer system **110**) through SPD device **124**.

FIG. 2 is a block diagram illustrating an implementation of data hardening module **126**, in accordance with some embodiments. Data hardening module **126** typically includes one or more processors (also sometimes called CPUs or processing units or microprocessors or microcontrollers) **202** for executing modules, programs and/or instructions stored in memory **206** and thereby performing processing operations, memory **206**, power storage and distribution module **250** (including energy storage device **204**), and one or more communication buses **208** for interconnecting these components. In some implementations, power storage and distribution module **250** includes circuitry for monitoring, storing, and distributing power for a storage device (e.g., storage device **120**, FIG. 1), including monitoring, controlling, charging, and/or testing energy storage device **204**. In some embodiments, energy storage device **204** includes one or more capacitors. In other embodiments, energy storage device **204** includes one or more inductors or any other passive elements that store energy.

Communication buses **208** optionally include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. Data harden-

ing module **126** is coupled to host interface **122**, SPD device **124**, memory controller **128**, and flash controllers **130** (e.g., flash controllers **130-1** through **130-m**) by communication buses **208**. Memory **206** includes high-speed random access memory, such as DRAM, SRAM, DDR RAM or other random access solid state memory devices, and may include non-volatile memory, such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid state storage devices. Memory **206** optionally includes one or more storage devices remotely located from processor(s) **202**. Memory **206**, or alternately the non-volatile memory device(s) within memory **206**, comprises a non-transitory computer readable storage medium. In some embodiments, memory **206**, or the computer readable storage medium of memory **206** stores the following programs, modules, and data structures, or a subset thereof:

- a monitor module **210** that is used for monitoring signals provided to a storage device (e.g., storage device **120**, FIG. **1**), for example to monitor and determine whether a power supply voltage provided to the storage device is lower than an under-voltage threshold;
- an energy storage device module **212** that is used for monitoring, controlling, charging, and/or testing an energy storage device (e.g., energy storage device **204**) on the storage device;
- a test module **214** that is used for testing one or more functions of the storage device;
- a power switch module **216** that is used for determining and controlling the voltage that is used to supply power to the storage device; and
- a power fail module **218** that is used for performing a power fail operation in accordance with a determination that a power supply voltage provided to the storage device is lower than an under-voltage threshold.

In some embodiments, memory **206**, or the computer readable storage medium of memory **206** further stores a configuration module for configuring storage device **120** and data hardening module **126**, and/or configuration values (such as one or more under-voltage threshold values) for configuring data hardening module **126**, neither of which is explicitly shown in FIG. **2**. In some implementations, upon power up and upon reset, the configuration module automatically sets the values of one or more configuration parameters of storage device **120** (and, optionally, determines which of two or more power fail modules, test modules, etc. to use) in accordance with the components of storage device **120** (e.g., the type of non-volatile memory components in storage device **120**) and/or characteristics of the data storage system **100** that includes storage device **120**.

In some embodiments, the power fail module **218** optionally includes the following modules or sub-modules, or a subset thereof:

- a signal module **220** that is used for signaling a power fail condition to a plurality of controllers on the storage device (e.g., memory controller **128** and flash controllers **130**, FIG. **1**);
- a reset module **222** that is used for resetting the plurality of controllers on the storage device;
- a power removal module **224** that is used for removing power from the plurality of controllers on the storage device;
- a guard banding module **226** that is used for tolerating power fluctuations in a power supply voltage provided to the storage device; and
- a discharge module **228** that is used for discharging the energy storage device on the storage device.

Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures or modules, and thus various subsets of these modules may be combined or otherwise re-arranged in various embodiments. In some embodiments, memory **206** may store a subset of the modules and data structures identified above. Furthermore, memory **206** may store additional modules and data structures not described above. In some embodiments, the programs, modules, and data structures stored in memory **206**, or the computer readable storage medium of memory **206**, provide instructions for implementing any of the methods described below with reference to FIGS. **6A-6E**.

Although FIG. **2** shows data hardening module **126**, FIG. **2** is intended more as a functional description of the various features which may be present in a data hardening module than as a structural schematic of the embodiments described herein. In practice, and as recognized by those of ordinary skill in the art, items shown separately could be combined and some items could be separated.

FIG. **3** is a block diagram illustrating an implementation of a memory controller **128**, in accordance with some embodiments. Memory controller **128** typically includes one or more processors (also sometimes called CPUs or processing units or microprocessors or microcontrollers) **302** for executing modules, programs and/or instructions stored in memory **306** and thereby performing processing operations, memory **306**, and one or more communication buses **308** for interconnecting these components. Communication buses **308** optionally include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. Memory controller **128** is coupled to host interface **122**, data hardening module **126**, and flash controllers **130** (e.g., flash controllers **130-1** through **130-m**) by communication buses **308**. Memory **306** includes high-speed random access memory, such as DRAM, SRAM, DDR RAM or other random access solid state memory devices, and may include non-volatile memory, such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid state storage devices. Memory **306** optionally includes one or more storage devices remotely located from processor(s) **302**. Memory **306**, or alternately the non-volatile memory device(s) within memory **306**, comprises a non-transitory computer readable storage medium. In some embodiments, memory **306**, or the computer readable storage medium of memory **306** stores the following programs, modules, and data structures, or a subset thereof:

- an interface module **310** that is used for communicating with other components, such as host interface **122**, data hardening module **126**, and flash controllers **130**;
- a reset module **312** that is used for resetting memory controller **128**; and
- a power fail module **314** that is used for performing a power fail operation.

In some embodiments, the power fail module **314** optionally includes a transfer module **316** that is used for transferring data held in volatile memory to non-volatile memory.

Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures or modules, and thus various

subsets of these modules may be combined or otherwise re-arranged in various embodiments. In some embodiments, memory 306 may store a subset of the modules and data structures identified above. Furthermore, memory 306 may store additional modules and data structures not described above. In some embodiments, the programs, modules, and data structures stored in memory 306, or the computer readable storage medium of memory 306, provide instructions for implementing respective operations in the methods described below with reference to FIGS. 6A-6E.

Although FIG. 3 shows a memory controller 128, FIG. 3 is intended more as a functional description of the various features which may be present in a memory controller than as a structural schematic of the embodiments described herein. In practice, and as recognized by those of ordinary skill in the art, items shown separately could be combined and some items could be separated.

FIG. 4 is a block diagram illustrating an implementation of a flash controller 130-1, in accordance with some embodiments. Flash controller 130-1 typically includes one or more processors (also sometimes called CPUs or processing units or microprocessors or microcontrollers) 402 for executing modules, programs and/or instructions stored in memory 306 and thereby performing processing operations, memory 406, and one or more communication buses 408 for interconnecting these components. Communication buses 408 optionally include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. Flash controller 130-1 is coupled to memory controller 128, data hardening module 126, and flash memory devices 140 (e.g., flash memory devices 140-1 through 140-n) by communication buses 408. Memory 406 includes high-speed random access memory, such as DRAM, SRAM, DDR RAM or other random access solid state memory devices, and may include non-volatile memory, such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid state storage devices. Memory 406 optionally includes one or more storage devices remotely located from processor(s) 402. Memory 406, or alternately the non-volatile memory device(s) within memory 406, comprises a non-transitory computer readable storage medium. In some embodiments, memory 406, or the computer readable storage medium of memory 406 stores the following programs, modules, and data structures, or a subset thereof:

- an interface module 410 that is used for communicating with other components, such as memory controller 128, data hardening module 126, and flash memory devices 140;
- a reset module 412 that is used for resetting flash controller 130-1; and
- a power fail module 414 that is used for performing a power fail operation.

In some embodiments, the power fail module 414 optionally includes a transfer module 416 that is used for transferring data held in volatile memory to non-volatile memory.

Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures or modules, and thus various subsets of these modules may be combined or otherwise re-arranged in various embodiments. In some embodiments, memory 406 may store a subset of the modules and data structures identified above. Furthermore, memory 406 may store additional modules and data structures not described

above. In some embodiments, the programs, modules, and data structures stored in memory 406, or the computer readable storage medium of memory 406, provide instructions for implementing respective operations in the methods described below with reference to FIGS. 6A-6E.

Although FIG. 4 shows a flash controller 130-1, FIG. 4 is intended more as a functional description of the various features which may be present in a flash controller than as a structural schematic of the embodiments described herein. In practice, and as recognized by those of ordinary skill in the art, items shown separately could be combined and some items could be separated. Further, although FIG. 4 shows a flash controller 130-1, the description of FIG. 4 similarly applies to other flash controllers (e.g., flash controllers 130-2 through 130-m) in storage device 120 (FIG. 1).

FIG. 5 is a block diagram illustrating an implementation of data hardening module 126, in accordance with some embodiments. While some example features are illustrated, various other features have not been illustrated for the sake of brevity and so as not to obscure more pertinent aspects of the example implementations disclosed herein. To that end, as a non-limiting example, data hardening module 126 includes processor 202, boost circuitry 520, hold-up capacitor 522, transistors 511-512, keeper circuitry 524, voltages V_{dd} 502, V_{SPD} 504, V_{holdup} 506, and $V_{switched}$ 508, and connections 530. In some implementations, V_{dd} 502 is a voltage supplied by a host system (e.g., computer system 110, FIG. 1) and has a target value of 1.5 volts or less. In some implementations, V_{holdup} 506 is a boosted up voltage from V_{dd} 502 and has a target value of 5.7 volts. In some embodiments, V_{holdup} 506 is used to charge an energy storage device (e.g., hold-up capacitor 522). In some implementations, V_{SPD} 504 is a voltage supplied for serial presence detect (SPD) functionality and has a target value of 3.3 volts. Further, in some implementations, only one of transistors 511, 512 is enabled at any one time. For example, whenever transistor 512 is enabled, transistor 511 is disabled (open state), so as to ensure that power from the data hardening circuit's energy storage device (e.g., hold-up capacitor 522) is not drained to the host system. Furthermore, whenever transistor 511 is enabled, providing power to components of memory device 120 from the host system, transistor 512 is disabled. In some implementations, the data hardening circuit's energy storage device (e.g., hold-up capacitor 522) stores, immediately prior to a power fail condition being detected, at least approximately 30 to 70 millijoules of energy per flash memory controller 130 in storage device 120.

In some implementations, processor 202 monitors and manages the functionality in data hardening module 126. For example, processor 202 monitors voltages V_{dd} 502 and V_{SPD} 504. If either V_{dd} 502 or V_{SPD} 504 fall below corresponding under-voltage thresholds, processor 202 signals a power fail condition to a plurality of controllers on storage device 120 (e.g., memory controller 128 and flash controllers 130, FIG. 1). In some embodiments, the under-voltage threshold varies depending on the target value of the voltage. For example, if the target voltage for V_{dd} 502 is 1.5 volts, the under-voltage threshold may be 1.5 volts minus 5% (i.e., 1.425 volts), so processor 202 would signal a power fail condition if V_{dd} 502 is lower than 1.425 volts. In some implementations, the under-voltage threshold for V_{dd} 502 is different than the under-voltage threshold for V_{SPD} 504.

In some embodiments, during regular operation of storage device 120, V_{dd} 502 is used to supply power to storage device 120. However, during a power fail operation, an energy storage device (e.g., hold-up capacitor 522) is used to provide power to storage device 120. In some implementations, pro-

11

cessor 202 controls transistors 511-512 to control $V_{switched}$ 508 to be voltage from V_{dd} 502 (e.g., during regular operation) or voltage from hold-up capacitor 522 (e.g., during a power fail operation). For example, during regular operation of storage device 120, V_{dd} 502 is used to supply power to storage device 120, so transistor 511 is turned on (e.g., to complete the connection between V_{dd} 502 and $V_{switched}$ 508) and transistor 512 is turned off (e.g., to disable the connection between hold-up capacitor 522 and $V_{switched}$ 508). However, during a power fail operation, hold-up capacitor 522 is used to provide power to storage device 120, so transistor 511 is turned off (e.g., to disable the connection between V_{dd} 502 and $V_{switched}$ 508) and transistor 512 is turned on (e.g., to enable the connection between hold-up capacitor 522 and $V_{switched}$ 508). Although a single hold-up capacitor 522 is shown in FIG. 5, any energy storage device, including one or more capacitors, one or more inductors, or one or more other passive elements that store energy, may be used to store energy to be used during a power fail operation.

In some implementations, hold-up capacitor 522 is charged using V_{holdup} 506, a voltage higher than V_{dd} 502. In some implementations, V_{dd} 502 is boosted up to V_{holdup} 506 using boost circuitry 520 (e.g., 1.35 volts or 1.5 volts is boosted up to 5.7 volts). In some implementations, boost circuitry 520 is controlled and enabled by processor 202. Further, in some embodiments, $V_{switched}$ 508 is used as an input to keeper circuitry 524, which along with V_{SPD} 504 provides power to processor 202. During a power fail operation, $V_{switched}$ 508 is provided via keeper circuitry 524 to processor 202 so as to provide power to processor 202 during the power fail operation. In some implementations, processor 202 has one or more connections 530 used to monitor and control other functions within storage device 120. In some implementations, V_{SPD} 504 provides power to keeper circuitry 524. Furthermore, in some implementations, V_{SPD} 504 is provided to storage device 120 before V_{dd} 502 is provided to storage device 120, allowing devices in storage device 120 to operate before main power V_{dd} 502 is provide to storage device 120.

FIGS. 6A-6E illustrate a flowchart representation of a method 600 of protecting data in a storage device, in accordance with some embodiments. A storage device (e.g., storage device 120, FIG. 1) coordinates and manages multiple sub-system components to protect data, which initiates performance of method 600. At least in some implementations, method 600 is performed by a storage device (e.g., storage device 120, FIG. 1) or one or more components of the storage device (e.g., data hardening module 126, memory controller 128, and/or flash controllers 130, FIG. 1). In some embodiments, method 600 is governed by instructions that are stored in a non-transitory computer readable storage medium and that are executed by one or more processors of a device, such as the one or more processors 202 of data hardening module 126, the one or more processors 302 of memory controller 128, and/or the one or more processors 402 of flash controllers 130, as shown in FIGS. 2-4.

A storage device (e.g., storage device 120, FIG. 1) determines (602) whether a power supply voltage provided to the storage device is lower than an under-voltage threshold. In some embodiments, the under-voltage threshold (sometimes also called a trip point) varies depending on the target value of the voltage. For example, if the target value of the power supply voltage is 1.5 volts, the under-voltage threshold may be 1.5 volts minus 5% (i.e., 1.425 volts), and the storage device determines whether the power supply voltage is lower than 1.425 volts. In some implementations, a monitor module (e.g., monitor module 210, FIG. 2) is used to determine

12

whether a power supply voltage provided to a storage device is lower than an under-voltage threshold, as described above with respect to FIG. 2.

In some embodiments, the power supply voltage is (604) a voltage supplied by a host system. In some implementations, the voltage supplied by a host system (e.g., V_{dd} 502, FIG. 5) has a target value of 1.5 volts or less. For example, for a double data rate type three (DDR3) interface specification, the supply voltage is 1.5 volts or 1.35 volts.

In some embodiments, the power supply voltage is (606) a voltage supplied for serial presence detect (SPD) functionality. In some implementations, the voltage supplied for SPD functionality (e.g., V_{SPD} 504, FIG. 5) has a target value of 3.3 volts.

In some embodiments, the storage device includes (608) a dual in-line memory module (DIMM) device. In some implementations, the storage device is compatible with a DIMM memory slot. For example, in some implementations, the storage device is compatible with a 240-pin DIMM memory slot using a DDR3 interface specification.

Next, the storage device, in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, performs (610) a power fail operation. Using the example above where the target value of the power supply voltage is 1.5 volts and the under-voltage threshold is 1.425 volts, in accordance with a determination that the power supply voltage is lower than 1.425 volts, the storage device performs a power fail operation. In some implementations, a power fail module (e.g., power fail module 218, FIG. 2) is used to perform a power fail operation, in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, as described above with respect to FIG. 2.

In some embodiments, one or more power supply voltages are monitored for under-voltage thresholds and the power fail operation is performed if any of the monitored power supply voltages fall below their respective under-voltage thresholds. For example, if two power supply voltages (e.g., a first power supply voltage and a second power supply voltages) are monitored for under-voltage thresholds (e.g., a first under-voltage threshold and a second under-voltage threshold, respectively), the power fail operation is performed in accordance with a determination that the first power supply voltage is lower than the first under-voltage threshold and the power fail operation is performed in accordance with a determination that the second power supply voltage is lower than the second under-voltage threshold. The under-voltage threshold (sometimes called "trip point") varies based on the target value of the power supply voltage. In some embodiments, the first power supply voltage is a voltage supplied by a host system (e.g., with a target value of 1.5 volts or less) and the second power supply voltage is a voltage supplied for serial presence detect (SPD) functionality (e.g., with a target value of 3.3 volts).

First, the power fail operation includes (610) signaling (612) a power fail condition to a plurality of controllers on the storage device (e.g., memory controller 128 and flash controllers 130, FIG. 1). In some implementations, a signal module (e.g., signal module 220, FIG. 2) is used to signal a power fail condition to a plurality of controllers on the storage device, as described above with respect to FIG. 2.

In some embodiments, the plurality of controllers on the storage device include (614) a memory controller (e.g., memory controller 128, FIG. 1) and one or more flash controllers (e.g., flash controllers 130, FIG. 1), the one or more

13

flash controllers coupled by the memory controller to a host interface (e.g., host interface **122**, FIG. **1**) of the storage device.

In some embodiments, the plurality of controllers on the storage device include **(616)** at least one non-volatile memory controller and at least one other memory controller other than the at least one non-volatile memory controller. In some implementations, the at least one non-volatile memory controller is a flash controller (e.g., flash controller **130-1**, FIG. **1**). In other implementations, the at least one non-volatile memory controller controls one or more other types of non-volatile memory devices.

In some embodiments, one of the plurality of controllers on the storage device maps **(618)** double data rate (DDR) interface commands to serial advance technology attachment (SATA) interface commands. For example, a memory controller (e.g., memory controller **128**, FIG. **1**) maps double data rate type three (DDR3) interface commands to SATA interface commands. In some implementations, a memory controller (e.g., memory controller **128**, FIG. **1**) uses a defined interface standard, such as DDR3, to communicate with a host interface (e.g., host interface **122**, FIG. **1**) and uses a defined interface standard, such as SATA, to communicate with other controllers on the storage device (e.g., flash controllers **130**, FIG. **1**).

Next, the power fail operation includes **(610)** transferring **(620)** data held in volatile memory to non-volatile memory (e.g., flash memory devices **140**, **142**, FIG. **1**). In some implementations, a power fail module on one or more controllers (e.g., power fail module **314**, FIG. **3** and power fail module **414**, FIG. **4**) are used to transfer data held in volatile memory to non-volatile memory, as described above with respect to FIGS. **3-4**.

In some embodiments, transferring data held in volatile memory to non-volatile memory includes transferring **(622)** data (e.g., volatile data **318**, FIG. **3**) from the memory controller (e.g., memory controller **128**, FIG. **1**) to the one or more flash controllers (e.g., flash controllers **130**, FIG. **1**). In some implementations, data transferred from the memory controller to the one or more flash controllers includes data in flight from the host interface (e.g., host interface **122**, FIG. **1**) to the memory controller, data that has been signaled to the host (e.g., computer system **110**, FIG. **1**) as saved (e.g., stored in a non-volatile store or write cache), and/or metadata stored in volatile memory in the memory controller. In some implementations, a transfer module (e.g., transfer module **316**, FIG. **3**) is used to transfer data from the memory controller to the one or more flash controllers, as described above with respect to FIG. **3**.

In some embodiments, transferring data held in volatile memory to non-volatile memory includes transferring **(624)** data (e.g., volatile data **418**, FIG. **4**) from the one or more flash controllers (e.g., flash controllers **130**, FIG. **1**) to the non-volatile memory (e.g., flash memory devices **140**, **142**, FIG. **1**). In some implementations, data transferred from the one or more flash controllers to the non-volatile memory includes data in flight to the one or more flash controllers and/or metadata stored in volatile memory in the one or more flash controllers (e.g., unwritten parity data; usage metrics representing or corresponding to cumulative usage and/or remaining endurance of respective portions (e.g., blocks, super-blocks, or other units) of the flash memory devices; address translation tables, etc.). In some implementations, a transfer module (e.g., transfer module **416**, FIG. **4**) is used to transfer data from the one or more flash controllers to the non-volatile memory, as described above with respect to FIG. **4**.

14

In some embodiments, the non-volatile memory comprises **(626)** one or more flash memory devices (e.g., flash memory devices **140**, **142**, FIG. **1**). In some implementations, the non-volatile memory includes a single flash memory device, while in other implementations the non-volatile memory includes a plurality of flash memory devices. In some implementations, the non-volatile memory includes NAND-type flash memory or NOR-type flash memory. In other embodiments, the non-volatile memory comprises one or more other types of non-volatile storage devices.

Next, the power fail operation includes **(610)** removing **(628)** power from the plurality of controllers on the storage device (e.g., memory controller **128** and flash controllers **130**, FIG. **1**). In some implementations, removing power from the plurality of controllers on the storage device include affirmatively removing power from the plurality of controllers (as opposed to allowing the plurality of controllers to automatically lose power). In some implementations, a power removal module (e.g., power removal module **224**, FIG. **2**) is used to remove power from the plurality of controllers on the storage device, as described above with respect to FIG. **2**.

In some embodiments, removing power from the plurality of controllers on the storage device includes resetting **(630)** the memory controller subsequent to transferring data from the memory controller to the one or more flash controllers. In some implementations, the memory controller (e.g., memory controller **128**, FIG. **1**) signals to a data hardening module (e.g., data hardening module **126**, FIG. **1**) when it has completed transferring data to the one or more flash controllers (e.g., flash controllers **130**, FIG. **1**). When the data hardening module receives this signal (e.g., a data hardening done signal), it can then reset the memory controller. In some implementations, resetting the memory controller puts the memory controller in a predefined low power state. In some implementations, resetting the memory controller is controlled by a data hardening module and reset functionality is disabled while data is being transferred to non-volatile memory (e.g., the data hardening module prevents the memory controller from being reset when the memory controller is transferring data held in volatile memory to non-volatile memory). In some implementations, a reset module in the data hardening module (e.g., reset module **222**, FIG. **2**) in conjunction with a reset module in the memory controller (e.g., reset module **312**, FIG. **3**) is used to reset the memory controller subsequent to transferring data from the memory controller to the one or more flash controllers, as described above with respect to FIGS. **2-3**.

Next, in some embodiments, removing power from the plurality of controllers on the storage device includes removing **(632)** power from the memory controller subsequent to resetting the memory controller. In some implementations, the memory controller (e.g., memory controller **128**, FIG. **1**) is in a separate power domain from the one or more flash controllers (e.g., flash controller **130**, FIG. **1**). Thus, power may be removed from the memory controller, regardless of whether the flash controllers are still transferring data to non-volatile memory. These independent power domains allow a data hardening module (e.g., data hardening module **126**, FIG. **1**) to selectively remove power from the various controllers on the storage device (e.g., storage device **120**, FIG. **1**). By selectively removing power to controllers as they are ready, the data hardening module preserves the power stored in the energy storage device (e.g., energy storage device **204**, FIG. **2**) for remaining tasks in the power fail operation. In some implementations, a power removal module (e.g., power removal module **224**, FIG. **2**) is used to

15

remove power from the memory controller subsequent to resetting the memory controller, as described above with respect to FIG. 2.

In some embodiments, the one or more flash controllers include (634) a first flash controller and a second flash controller and removing power from the plurality of controllers on the storage device includes resetting (636) the first flash controller subsequent to transferring data from the first flash controller to the non-volatile memory. In some implementations, the first flash controller (e.g., flash controller 130-1, FIG. 1) signals to a data hardening module (e.g., data hardening module 126, FIG. 1) when it has completed transferring data to the non-volatile memory (e.g., flash memory devices 140). When the data hardening module receives this signal (e.g., a data hardening done signal), it can then reset the first flash controller. In some implementations, resetting the first flash controller puts the first flash controller in a predefined low power state. In some implementations, the data hardening module prevents the first flash controller from being reset when the first flash controller is transferring data held in volatile memory to non-volatile memory. In some implementations, a reset module in the data hardening module (e.g., reset module 222, FIG. 2) in conjunction with a reset module in the first flash controller (e.g., reset module 412, FIG. 4) is used to reset the first flash controller subsequent to transferring data from the first flash controller to the non-volatile memory, as described above with respect to FIGS. 2 and 4.

In some embodiments, the one or more flash controllers include (634) a first flash controller and a second flash controller and removing power from the plurality of controllers on the storage device further includes resetting (638) the second flash controller subsequent to transferring data from the second flash controller to the non-volatile memory. Explanations provided above in connection with resetting the first flash controller (in operation 636) are equally applicable to resetting the second flash controller. In some implementations, a reset module in the data hardening module (e.g., reset module 222, FIG. 2) in conjunction with a reset module in the second flash controller (e.g., reset module 412, FIG. 4) is used to reset the second flash controller subsequent to transferring data from the second flash controller to the non-volatile memory, as described above with respect to FIGS. 2 and 4.

In some embodiments, the one or more flash controllers include (634) a first flash controller and a second flash controller and removing power from the plurality of controllers on the storage device further includes removing (640) power from the first and the second flash controllers subsequent to resetting the first and second flash controllers. In some embodiments, the first flash controller and the second flash controller share the same power domain, and power is removed from the first and the second flash controllers after both the first and the second flash controllers have been reset. In some embodiments, the first flash controller is in a first power domain and the second flash controller is in a second power domain, and power is removed from the first flash controller independent of when power is removed from the second flash controller. In some implementations, a power removal module (e.g., power removal module 224, FIG. 2) is used to remove power from the first and the second flash controllers subsequent to resetting the first and second flash controllers, as described above with respect to FIG. 2.

In some embodiments, removing (642) power from the first and the second flash controllers is subsequent to removing power from the memory controller. As discussed above, independent power domains on the storage device allow a data

16

the storage device (e.g., memory controller 128 and flash controllers 130, FIG. 1). In some implementations, data hardening is completed in a cascading manner since the flash controllers cannot finish hardening data until they have received all the volatile data that needs to be transferred from the memory controller. As a result, the memory controller will complete its data hardening first and thus, power will be removed from the memory controller before power is removed from the first and the second flash controllers.

In some embodiments, the power fail operation is (644) performed to completion regardless of whether the power supply voltage returns to a voltage greater than or equal to the under-voltage threshold. In some implementations, even if the power fail condition is temporary (e.g., a lightning strike that briefly causes the power supply voltage to flicker below the under-voltage threshold), the power fail operation is performed to completion. In some implementations, once a power fail operation begins, a data hardening module (e.g., data hardening module 126, FIG. 1) effectively disconnects from the power supply voltage provided to the storage device and ignores the power supply voltage until the power fail operation is complete.

In some embodiments, the power supply voltage includes (646) a first voltage and a second voltage, and performing the power fail operation includes (610) performing (648) the power fail operation in accordance with a determination that the first voltage is lower than a first under-voltage threshold. In some embodiments, the first voltage is a voltage supplied by a host system (e.g., V_{dd} 502, FIG. 5) and the first voltage has a target value of 1.5 volts or less. For example, for a double data rate type three (DDR3) interface specification, the voltage supplied by a host system is 1.5 volts or 1.35 volts.

Next, where the power supply voltage includes (646) a first voltage and a second voltage, performing the power fail operation includes (610) performing (650) the power fail operation in accordance with a determination that the second voltage is lower than a second under-threshold voltage. In some embodiments, the second voltage is a voltage supplied for serial presence detect (SPD) functionality (e.g., V_{SPD} 504, FIG. 5) and the second voltage has a target value of 3.3 volts.

In some embodiments, the power fail operation is (652) performed using power from an energy storage device (e.g., energy storage device 204, FIG. 2) on the storage device (e.g., storage device 120, FIG. 1). As described above with respect to FIG. 5, during a power fail operation, an energy storage device (e.g., hold-up capacitor 522, FIG. 5) is used to provide power to the storage device, and a data hardening module (e.g., data hardening module 126, FIGS. 1 and 5) is used to connect and disconnect the appropriate power sources (e.g., disabling the connection between V_{dd} 502 and $V_{switched}$ 508 and enabling the connection between hold-up capacitor 522 and $V_{switched}$ 508, FIG. 5). Alternatively, in some embodiments, the power fail operation is performed using power from a reserve energy storage device that is located external to the storage device (e.g., storage device 120, FIG. 1).

In some embodiments, the energy storage device includes (654) one or more capacitors. For example, in some implementations, the energy storage device includes a single capacitor (e.g., hold-up capacitor 522, FIG. 5), while in other implementations, the energy storage device includes a plurality of capacitors. In other implementations, the energy storage device includes one or more inductors. In some implementations, the energy storage device includes one or more other passive elements that store energy.

Optionally, the storage device monitors (656) the energy storage device to ensure capacitors in the energy storage device are charged to at least a first charge level. In some

implementations, a data hardening module (e.g., data hardening module **126**, FIG. **1**) performs health monitoring for the energy storage device and monitors the capacitor(s) in the energy storage device. For example, if the first charge level is 98% (or, alternatively, 95% or 90%) of charge capacity, the data hardening module monitors the energy storage device to ensure the capacitor(s) in the energy storage device are charged to at least 98% (or, alternatively, 95% or 90%) of charge capacity. In some implementations, if one or more capacitors of the energy storage device are not charged to at least the first charge level, the storage device performs one or more remedial actions (e.g., further charging the one or more capacitors). In some implementations, an energy storage device module (e.g., energy storage device module **212**, FIG. **2**) is used to monitor the energy storage device to ensure capacitors in the energy storage device are charged to at least a first charge level, as described above with respect to FIG. **2**.

Further, the storage device selectively tests (**658**) one or more capacitors from the energy storage device during operation of the storage device. In some implementations, the data hardening module (e.g., data hardening module **126**, FIG. **1**) uses an algorithm to selectively test one or more capacitors from the energy storage device during operation of the storage device. In some implementations, one or more capacitors from the energy storage device are tested during regular operation of the storage device (as opposed to during a power fail operation). Since testing the capacitor(s) will discharge the capacitor(s), the data hardening module manages the coordination of testing the capacitor(s) to ensure that testing of the capacitor(s) from the energy storage device does not interfere with a potential power fail operation. In some implementations, an energy storage device module (e.g., energy storage device module **212**, FIG. **2**) is used to selectively test one or more capacitors from the energy storage device during operation of the storage device, as described above with respect to FIG. **2**.

Optionally, prior to determining (**660**) whether the power supply voltage provided to the storage device is lower than the under-voltage threshold, the storage device charges (**662**) the energy storage device using a higher voltage than the power supply voltage provided to the storage device. As described above with respect to FIG. **5**, an energy storage device (e.g., hold-up capacitor **522**, FIG. **5**) is charged using a higher voltage (e.g., V_{holdup} **506**, FIG. **5**) than the power supply voltage provided to the storage device (e.g., V_{dd} **502** or V_{SPD} **504**, FIG. **5**). For example, in some implementations, the energy storage device is charged using V_{holdup} (e.g., V_{holdup} **506**, FIG. **5**, with a target value of 5.7 volts), which is a higher voltage than V_{dd} **502** or V_{SPD} **504**, FIG. **5**, with target values of 1.35/1.5 volts and 3.3 volts, respectively). In some implementations, the energy storage device is charged with dynamic inrush current control. In some implementations, an energy storage device module (e.g., energy storage device module **212**, FIG. **2**) is used to charge the energy storage device using a higher voltage than the power supply voltage provided to the storage device, as described above with respect to FIG. **2**.

Next, the storage device determines (**664**) whether the energy storage device (e.g., energy storage device **204**, FIG. **2**) meets a minimum charge level threshold within a predefined charge time. For example, if the minimum charge level threshold is 95% charge capacity and the predefined charge time is 25 ms, the storage device determines whether the energy storage device is charged to at least 95% charge capacity within 25 ms. In some implementations, an energy storage device module (e.g., energy storage device module **212**, FIG. **2**) is used to determine whether the energy storage

device meets a minimum charge level threshold within a predefined charge time, as described above with respect to FIG. **2**.

Further, in accordance with a determination that the energy storage device does not meet the minimum charge level threshold in the predefined charge time, the storage device prevents (**666**) operation of the storage device. In some implementations, a determination that the energy storage device does not meet the minimum charge level threshold in the predefined charge time indicates that there will be a data hardening failure when a power fail operation is performed in the future (e.g., a predictive failure detection). As a result, operation of the storage device is prevented to avoid a future data hardening failure. In some implementations, an energy storage device module (e.g., energy storage device module **212**, FIG. **2**) is used to prevent operation of the storage device, in accordance with a determination that the energy storage device does not meet the minimum charge level threshold in the predefined charge time, as described above with respect to FIG. **2**.

In some embodiments, preventing operation of the storage device includes communicating (**668**) a failure message to a host system (e.g., computer system **110**, FIG. **1**). In some implementations, the failure message is communicated with the host system through a SPD device (e.g., SPD device **124**, FIG. **1**) or a SPD bus controller.

Optionally, the storage device discharges (**670**) the energy storage device subsequent to removing power from the plurality of controllers on the storage device. In some implementations, subsequent to removing power from the plurality of controllers (e.g., memory controller **128** and flash controllers **130**, FIG. **1**) on the storage device (e.g., storage device **120**, FIG. **1**), the storage device discharges the energy storage device (e.g., energy storage device, FIG. **2**) and the storage device shuts down like a regular hard power down. In some implementations, discharging the energy storage device includes discharging one or more capacitors of the energy storage device. In some implementations, a discharge module (e.g., discharge module **228**, FIG. **2**) is used to discharge the energy storage device subsequent to removing power from the plurality of controllers on the storage device, as described above with respect to FIG. **2**.

In some implementations, with respect to any of the methods described above, the non-volatile memory is a single flash memory device, while in other implementations, the non-volatile memory includes a plurality of flash memory devices.

In some implementations, with respect to any of the methods described above, a storage device includes (1) an interface for coupling the storage device to a host system, (2) a plurality of controllers, each of the plurality of controllers configured to transfer data held in volatile memory to non-volatile memory, and (3) a data hardening module including one or more processors and an energy storage device, the storage device configured to perform or control performance of any of the methods described above.

It will be understood that, although the terms “first,” “second,” etc. may be used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element from another. For example, a first contact could be termed a second contact, and, similarly, a second contact could be termed a first contact, which changing the meaning of the description, so long as all occurrences of the “first contact” are renamed consistently and all occurrences of the second contact are renamed consistently. The first contact and the second contact are both contacts, but they are not the same contact.

19

The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the claims. As used in the description of the embodiments and the appended claims, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term “and/or” as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms “comprises” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

As used herein, the term “if” may be construed to mean “when” or “upon” or “in response to determining” or “in accordance with a determination” or “in response to detecting,” that a stated condition precedent is true, depending on the context. Similarly, the phrase “if it is determined [that a stated condition precedent is true]” or “if [a stated condition precedent is true]” or “when [a stated condition precedent is true]” may be construed to mean “upon determining” or “in response to determining” or “in accordance with a determination” or “upon detecting” or “in response to detecting” that the stated condition precedent is true, depending on the context.

The foregoing description, for purpose of explanation, has been described with reference to specific implementations. However, the illustrative discussions above are not intended to be exhaustive or to limit the claims to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The implementations were chosen and described in order to best explain principles of operation and practical applications, to thereby enable others skilled in the art.

What is claimed is:

1. A method of protecting data in a storage device, the method comprising:

determining whether a power supply voltage provided to the storage device is lower than an under-voltage threshold; and

in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, performing a power fail operation, the power fail operation including:

signaling a power fail condition to a plurality of controllers on the storage device, the plurality of controllers including a memory controller and one or more flash controllers;

transferring data held in volatile memory to non-volatile memory, including:

transferring data, including metadata stored in volatile memory of the memory controller, from the memory controller to the one or more flash controllers; and

transferring data, including the metadata, from the one or more flash controllers to the non-volatile memory; and

removing power from the plurality of controllers on the storage device.

2. The method of claim 1, wherein the one or more flash controllers are coupled by the memory controller to a host interface of the storage device.

3. The method of claim 1, wherein removing power from the plurality of controllers on the storage device includes:

20

resetting the memory controller subsequent to transferring data from the memory controller to the one or more flash controllers; and

removing power from the memory controller subsequent to resetting the memory controller.

4. The method of claim 1, wherein the one or more flash controllers include a first flash controller and a second flash controller, and wherein removing power from the plurality of controllers on the storage device includes:

resetting the first flash controller subsequent to transferring data from the first flash controller to the non-volatile memory;

resetting the second flash controller subsequent to transferring data from the second flash controller to the non-volatile memory; and

removing power from the first and the second flash controllers subsequent to resetting the first and second flash controllers.

5. The method of claim 4, wherein removing power from the first and the second flash controllers is subsequent to removing power from the memory controller.

6. The method of claim 1, wherein the power fail operation is performed to completion regardless of whether the power supply voltage returns to a voltage greater than or equal to the under-voltage threshold.

7. The method of claim 1, wherein the power supply voltage is a voltage supplied by a host system.

8. The method of claim 1, wherein the power supply voltage is a voltage supplied for serial presence detect (SPD) functionality.

9. The method of claim 1, wherein the power supply voltage includes a first voltage and a second voltage, and wherein performing the power fail operation includes:

performing the power fail operation in accordance with a determination that the first voltage is lower than a first under-voltage threshold; and

performing the power fail operation in accordance with a determination that the second voltage is lower than a second under-voltage threshold.

10. The method of claim 1, wherein the power fail operation is performed using power from a reserve energy storage device.

11. The method of claim 1, wherein the power fail operation is performed using power from an energy storage device on the storage device.

12. The method of claim 11, wherein the energy storage device includes one or more capacitors.

13. The method of claim 11, further comprising: monitoring the energy storage device to ensure capacitors in the energy storage device are charged to at least a first charge level; and

selectively testing one or more capacitors from the energy storage device during operation of the storage device.

14. The method of claim 11, further comprising: prior to determining whether the power supply voltage provided to the storage device is lower than the under-voltage threshold:

charging the energy storage device using a higher voltage than the power supply voltage provided to the storage device;

determining whether the energy storage device meets a minimum charge level threshold within a predefined charge time; and

in accordance with a determination that the energy storage device does not meet the minimum charge level threshold in the predefined charge time, preventing operation of the storage device.

21

15. The method of claim 14, wherein preventing operation of the storage device includes communicating a failure message to a host system.

16. The method of claim 11, further comprising discharging the energy storage device subsequent to removing power from the plurality of controllers on the storage device. 5

17. The method of claim 1, wherein the non-volatile memory comprises one or more flash memory devices.

18. The method of claim 1, wherein the storage device includes a dual in-line memory module (DIMM) device. 10

19. The method of claim 1, wherein the plurality of controllers on the storage device include at least one non-volatile memory controller and at least one other memory controller other than the at least one non-volatile memory controller. 15

20. The method of claim 1, wherein one of the plurality of controllers on the storage device maps double data rate (DDR) interface commands to serial advance technology attachment (SATA) interface commands.

21. A storage device, comprising:

an interface for coupling the storage device to a host system;

a plurality of controllers including a memory controller and one or more flash controllers, each of the plurality of controllers configured to transfer data held in volatile memory to non-volatile memory; and 25

a data hardening module including one or more processors and an energy storage device, the data hardening module configured to:

determine whether a power supply voltage provided to the storage device is lower than an under-voltage threshold; and 30

in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, perform a power fail operation, the power fail operation including: 35

signaling a power fail condition to the plurality of controllers, causing the plurality of controllers to transfer data held in volatile memory to non-volatile memory, the transfer including:

transferring data, including metadata stored in volatile memory of the memory controller, from the memory controller to the one or more flash controllers; and 40

22

transferring data, including the metadata, from the one or more flash controllers to the non-volatile memory; and

removing power from the plurality of controllers on the storage device.

22. The storage device of claim 21, wherein the one or more flash controllers are coupled by the memory controller to a host interface of the storage device.

23. A non-transitory computer readable storage medium, storing one or more programs for execution by one or more processors of a storage device having a plurality of controllers including a memory controller and one or more flash controllers and a data hardening module, the one or more programs including instructions that when executed by the one or more processors of the storage device cause the data hardening module to: 15

determine whether a power supply voltage provided to the storage device is lower than an under-voltage threshold; and

in accordance with a determination that the power supply voltage is lower than the under-voltage threshold, perform a power fail operation, the power fail operation including:

signaling a power fail condition to the plurality of controllers, causing the plurality of controllers to transfer data held in volatile memory to non-volatile memory, the transfer including:

transferring data, including metadata stored in volatile memory of the memory controller, from the memory controller to the one or more flash controllers; and

transferring data, including the metadata, from the one or more flash controllers to the non-volatile memory; and

removing power from the plurality of controllers on the storage device.

24. The non-transitory computer readable storage medium of claim 23, wherein the non-transitory computer readable storage medium includes a non-transitory computer readable storage medium associated with each of the plurality of controllers on the storage device and a non-transitory computer readable storage medium associated with the data hardening module.

* * * * *